

# Knowing What’s Missing: Assessing Information Sufficiency in Question Answering

Akriti Jain, Aparna Garimella  
Adobe Research, India  
{akritij, garimell}@adobe.com

## Abstract

Determining whether a provided context contains sufficient information to answer a question is a critical challenge for building reliable question-answering systems. While simple prompting strategies have shown success on factual questions, they frequently fail on inferential ones that require reasoning beyond direct text extraction. We hypothesize that asking a model to first reason about what specific information is missing provides a more reliable, implicit signal for assessing overall sufficiency. To this end, we propose a structured *Identify-then-Verify* framework for robust sufficiency modeling. Our method first generates multiple hypotheses about missing information and establishes a semantic consensus. It then performs a critical verification step, forcing the model to re-examine the source text to confirm whether this information is truly absent. We evaluate our method against established baselines across diverse multi-hop and factual QA datasets. The results demonstrate that by guiding the model to justify its claims about missing information, our framework produces more accurate sufficiency judgments while clearly articulating any information gaps.

## 1 Introduction

Large Language Models (LLMs) have demonstrated remarkable capabilities in question answering (QA) (Ma et al., 2025; Lauriola et al., 2025; Shi et al., 2025; Yue, 2025; Fischer et al., 2024). However, a critical limitation undermines their reliability: these models often generate confident-sounding but unfaithful or hallucinated responses when the provided context lacks sufficient information to answer a question (Fadeeva et al., 2025; Ming et al., 2025; Joren et al., 2025; Krishna and et al., 2024; Ji et al., 2023).

Recent efforts to address this challenge have primarily relied on simple prompt-based strategies that directly ask models to assess whether sufficient information is available (Joren et al., 2025).

Such approaches show promise for factual questions, both single-hop (e.g., “What is the capital of France?”) and multi-hop (e.g., “Who directed the movie that won Best Picture in 1994?”), where the task reduces to verifying the presence of specific entities or factual chains. However, they frequently fail on inferential questions that require synthesizing information not directly stated in the text (e.g., “Why did the peace negotiations fail according to the sources?”) as illustrated in Figure 1. While models correctly identify missing factual links in multi-hop chains, they struggle to assess sufficiency for questions requiring inference. The core issue is that verifying keyword presence differs fundamentally from evaluating whether available facts provide adequate logical foundation for reasoning.

Existing prompting methods typically fall into two categories. *Direct sufficiency prompting* asks for an explicit binary judgment (sufficient or insufficient). While effective for fact-checking, this approach proves unreliable when sufficiency depends on whether scattered information collectively supports an inference. *Answering-based methods* (Wang et al., 2024) infer sufficiency from whether a model provides an answer versus abstaining. However, models often synthesize plausible-sounding answers from inadequate context using parametric knowledge, or hallucinate justifications for abstention that are not grounded in the text.

We hypothesize that a more robust approach for determining context sufficiency is to frame the task around identifying what specific information is missing. This reformulation offers two key advantages. First, it forces the model to form a concrete hypothesis about an information gap, which can lead to a more reasoned sufficiency judgment. Second, it produces a granular and actionable output that can be highly beneficial for downstream applications. For example, guiding the next query in an iterative retrieval system to find the necessary context can even improve the QA performance. Recent

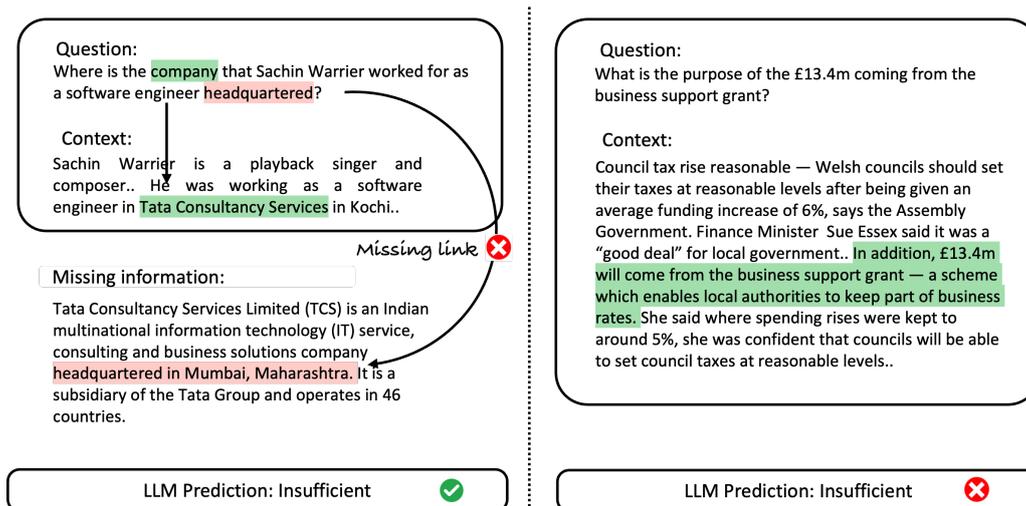


Figure 1: **LLMs Excel at Factual Verification but Falter on Inferential Reasoning.** (Left) For a fact-based multi-hop question, the model correctly identifies a missing factual link (the location of TCS headquarters) and accurately classifies the context as *Insufficient*. (Right) Conversely, for a question requiring simple inference, the model fails. It incorrectly predicts *Insufficient* when asked for the purpose of a business grant, unable to deduce that the scheme’s stated function, a scheme which enables local authorities to keep part of business rates-directly serves as its purpose. This highlights a key limitation where models struggle to assess sufficiency beyond literal fact-checking.

work on utility-based passage selection has shown that modeling which passages actually contribute to answering a question can significantly improve both retrieval and downstream performance (Zhang et al., 2025; Jain and Garimella, 2025; Xu et al., 2025). Our work complements this line of research by focusing on whether the retrieved context is collectively sufficient, rather than scoring individual passages. However, simply asking for missing information may not suffice, as the model may hallucinate non-existent gaps. To address this, we propose an **Identify-then-Verify** framework that first generates multiple hypotheses about missing information and establishes a consensus, then verifies this consensus against the source text. If the predicted information is indeed missing, the second step would fail to attribute it back, and we use this signal as the verification for the first stage’s output.

We make three main contributions in this paper. (1) We establish that information sufficiency prediction remains a challenging problem for open-ended inferential questions by benchmarking existing approaches on various multi-hop and answerability datasets. (2) We propose an Identify-then-Verify framework that reformulates sufficiency assessment as identifying and verifying missing information, leveraging self-consistency and semantic consensus to mitigate hallucination. Our approach can comprehensively identify information

gaps, providing actionable outputs for downstream applications. (3) We demonstrate substantial improvements over strong baselines across diverse benchmarks on complex questions both in accuracy and justification alignment. We further establish that sufficiency exists on a spectrum from pragmatic inference to strict literal matching, which our framework accommodates by treating verification strictness as a tunable parameter.

## 2 Related Work and Background

### 2.1 Information Sufficiency and Unanswerable Questions

For Large Language Models (LLMs) to be truly reliable, they must not only provide accurate answers but also recognize when they lack the information to do so. The study of unanswerable questions has become a widely used framework to investigate this behavior, probing issues such as model hallucination (Kim et al., 2024; Tomani et al., 2024), context faithfulness (Ming et al., 2025), and the fundamental question of whether LLMs know when not to answer (Saadat et al., 2024; Madhusudhan et al., 2024; Deng et al., 2024).

Successful abstention requires models to reason about both the evidence provided and the uncertainty of their conclusions (Kirichenko et al., 2025). However, a critical distinction must be made: is

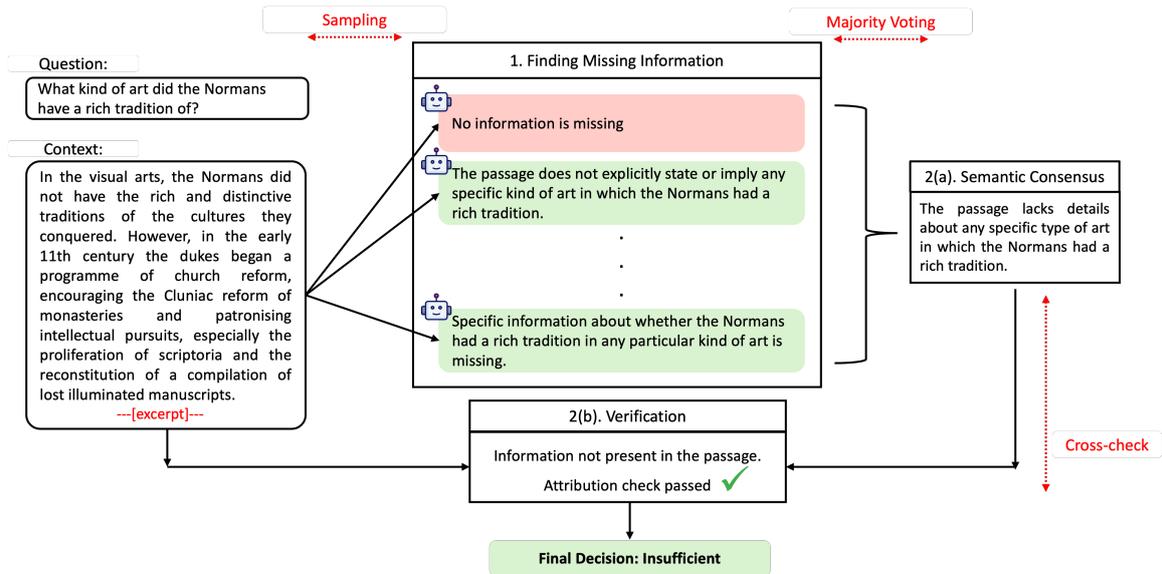


Figure 2: **An Overview of Our Identify-then-Verify Framework.** Given a question and context, the framework first queries an LLM multiple times with non-zero temperature to generate a diverse set of hypotheses about missing information (**Step 1: Identify**). This captures a rich distribution of the model’s reasoning. Next, a single Consensus Gap Claim is established from these hypotheses (**Step 2a**). This claim is then checked against the original context in a final verification step (**Step 2b**). This verification acts as a critical self-correction mechanism, with the final sufficiency decision based on its outcome.

the model unable to utilize the context correctly, or is the context itself insufficient? Joren et al. (2025) formalize this distinction by introducing the concept of *sufficient context* and demonstrate that LLMs can directly classify sufficiency for many cases. Our work investigates the limitations of such direct classification approaches, particularly for complex inferential questions, and proposes an alternative formulation based on explicit gap identification.

## 2.2 Missing Information Detection

The concept of prompting a model to identify missing information has been explored for guiding information retrieval as seen in frameworks like MIGRES (Wang et al., 2024). However, the reliability of this method remains a challenge. As demonstrated by benchmarks like AbsenceBench (Fu et al., 2025), while models excel at recalling surprising information, they still struggle to identify clearly omitted information. On the other hand, recent work on Missing Premise (MiP) problems reveals another failure mode: MiP-Overthinking (Fan et al., 2025). The study finds that even when reasoning models detect a missing premise early in their thought process, they often fail to commit to this judgment. Instead of abstaining, they fall into excessively redundant and meaningless thinking,

generating significantly longer responses while repeatedly attempting to solve an unsolvable problem. Taken together, these findings show that a single, greedy generation from an LLM is an unstable and insufficient signal. The model might fail to identify the gap entirely, or it might notice the gap but get stuck in a noisy, ineffective reasoning loop.

## 2.3 Context Evaluation vs. Context Compression

A related but distinct line of research focuses on optimizing computational efficiency in context processing. The work of (Xie et al., 2025) on dynamic context cutoff, for example, seeks to reduce token processing by identifying the earliest point at which an LLM has enough information to answer, effectively performing on-the-fly compression by analyzing internal model states (e.g., attention activations) to discard redundant text.

Our framework differs fundamentally in both goal and approach. Rather than optimizing for efficiency through compression, we perform *evaluation* of context completeness for reliable answering. We analyze the entire, unaltered input to determine whether it contains sufficient information, asking not whether the context can be made smaller, but whether it is adequate in its original form. These approaches are complementary: context compres-

sion assumes sufficiency and optimizes processing, while our method explicitly assesses sufficiency to prevent unreliable outputs.

### 3 Problem Formulation and Methodology

We hypothesize that prompting a model for a simple binary prediction (Sufficient/Insufficient) is sub-optimal. This approach is brittle, particularly for complex inferential questions where the answer does not exist in specific spans but must be inferred or synthesized from multiple portions of the context. A model might correctly classify context as insufficient but for an incorrect reason, a nuance lost in binary labels and poorly captured by standard accuracy metrics. We therefore reformulate the task from a classification problem to a generative one: *What specific information is missing?* This shift forces the model to produce a falsifiable claim, offering a more transparent and interpretable signal of its reasoning. We propose a structured, two-step Identify-then-Verify framework designed to leverage the generative power of LLMs while mitigating their risks, such as hallucination and inconsistency (Fig. 2).

#### Step 1: Identify Potential Gaps via Self-Consistency.

Rather than asking a model to predict whether context is sufficient, we prompt it to identify *what* information is missing to answer the question. However, LLM responses for missing information detection are brittle, as documented by Fu et al. (2025). To quantify this instability, we analyze hypothesis disagreement (Figure 3): the frequency with which a model produces conflicting judgments about the same question-context pair across multiple runs. Our analysis reveals two key findings. First, disagreement increases sharply from one to two runs, providing strong evidence that single-run generation is unreliable. Table 1 illustrates such disagreement. Second, the disagreement rate stabilizes after 4 runs. Based on these findings, we query the model  $N = 5$  times with non-zero temperature to generate diverse hypotheses about missing information.

This approach is motivated by recent work demonstrating that aggregating multiple LLM outputs significantly improves reliability and provides robust uncertainty quantification (Wang et al., 2023; Yang et al., 2025; Li et al., 2025). Sampling with non-zero temperature encourages exploration of different plausible token sequences, yielding a

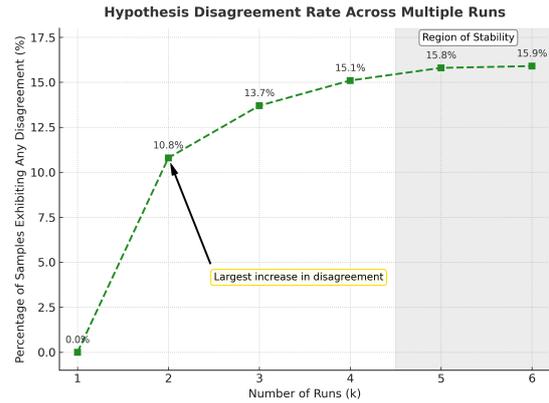


Figure 3: **Rate of Hypothesis Disagreement Across Multiple Runs.** The plot shows the percentage of samples with conflicting judgments for a given number of runs. The sharpest rise occurs when moving from one to two runs, highlighting the risk of a single pass. The curve flattens significantly after four runs, motivating our choice of  $N = 5$  as the point of stability.

distribution of potential information gaps rather than a single deterministic output. This distribution provides a richer signal reflecting the model’s assessment of ambiguity.

**Step 2: Establish Consensus and Verify.** The output of the Identify step is a set of natural language statements, and not classification labels. Naive majority voting is therefore inapplicable, as different phrasings may convey the same semantic meaning (e.g., “the year isn’t mentioned,” “it lacks a specific date,” “the date is missing”).

**Semantic Consensus.** We require a more sophisticated aggregation method than a simple majority vote. To facilitate this, we first encode all gap-identifying hypotheses into dense vector embeddings using a sentence Transformer model. From these embeddings, we compute a matrix of cosine similarities. The hypothesis with the highest average similarity to all others is selected as the *Consensus Gap Claim*. This ensures we capture the semantic “center of gravity” of the model’s collective reasoning. If no single idea achieves a minimum semantic similarity threshold, we recognize that the model is highly uncertain and can conservatively default to an insufficient judgment.

**Attribution Verification.** This Consensus Gap Claim is then subjected to a final, critical attribution check. We initiate a second, independent LLM call with a focused and less ambiguous task: verify if the information specified in the consensus claim is, in fact, absent from the source passage(s). This verification step acts as a self-correction mecha-

<b>Question:</b>	How much did Metallica earn from their 2004 tour?
<b>Key Context:</b>	Heavy metal band Metallica saw their 2004 earnings rise to \$43.1 million (£23.1 million), thanks to their <i>Madly in Anger with the World</i> tour. [Context excerpt from music industry earnings report...]
<b>Ground Truth:</b>	<b>Sufficient</b>
<b>Prediction 1:</b>	<b>Insufficient</b>
	<i>Reasoning:</i> The context does not provide sufficient information to directly answer the question. While it mentions Metallica’s total earnings in 2004, it does not specify how much came solely from their tour.
<b>Prediction 2:</b>	<b>Sufficient</b>
	<i>Reasoning:</i> The context provides all necessary information to answer the question. It explicitly states that Metallica earned \$43.1m from their 2004 tour.

Table 1: Example of sufficiency disagreement reasoning across runs.

nism. By decoupling the exploratory *identification* step from the factual *verification* step, we prevent the model from finalizing a decision based on a hallucinated or misremembered information gap. The final sufficiency determination is based on this attribution check: if the verifier finds the claimed missing information in the context, we classify the context as Sufficient; otherwise, Insufficient.

## 4 Experiments

In this section, we discuss the details of our experimental setup, datasets used, baselines, and evaluation measures.

### 4.1 Datasets

We evaluate our framework on two categories of datasets, namely multi-hop question answering and answerability benchmarks.

**MultihopQA datasets:** Since these datasets contain questions that are inherently answerable, we create negative (insufficient) samples by removing the designated supporting passages from the context. **HotpotQA** (Yang et al., 2018) features questions that require 2-hop reasoning. Each question is accompanied by 2 gold supporting passages and 8 distractor passages. **MuSiQue** (Trivedi et al., 2022) is designed for compositional reasoning. The questions involve 2 to 4 reasoning hops with 20 total passages. We use two settings: MuSiQue-Ans, containing only answerable questions, and MuSiQue-Full, a more stringent and hard version with unanswerable contrast questions designed to prevent disconnected reasoning. **2WikiMultiHopQA** (Ho et al., 2020) contains 2 or 4-hop questions con-

structed from Wikipedia, requiring the model to combine information from two different Wikipedia entities or facts to arrive at the answer. Each question is provided with 10 passages.

**Answerability benchmarks:** These datasets are specifically designed to test models’ ability to identify unanswerable questions.

**CouldAsk** (Zhao et al., 2024) studies how models handle questions with presupposition errors, where assumptions conflict with or are unsupported by the context. It includes existing datasets (SQuADv2 (Rajpurkar et al., 2018), QA2 (Kim et al., 2023), BanditQA (Gao et al., 2022)) and three synthetically generated datasets (BBC, Reddit, Yelp) designed to confuse LLMs across diverse domains beyond Wikipedia. The overall performance on these newer datasets (BBC, Reddit, Yelp) for detecting unanswerable questions is often lower than on existing ones, confirming their challenging nature. **FaithEval** (Ming et al., 2025) evaluates the contextual faithfulness of LLMs. We focus on its unanswerable setting, where the context is relevant but lacks the specific information needed to answer.

**Data Filtering.** Following Dai et al. (2025), we filter out questions that LLMs can answer correctly without any context to mitigate parametric knowledge effects. Such cases would inflate sufficiency detection performance, as models could deem insufficient context as “sufficient” by filling gaps from internal memory. Our filtering procedure is as follows: for each question in a given dataset, we first prompt an LLM (GPT-4o) to generate an answer without providing any context. This context-free answer is then evaluated for factual correctness using a second LLM as judge. The evaluator is presented with the original question, the ground-truth answer, and the model’s context-free answer, and is tasked with determining if the generated answer is CORRECT, INCORRECT, or UNCLEAR. Any question for which the model’s context-free answer is judged as CORRECT is considered “known” and is subsequently removed from the evaluation set. In cases where the ground-truth answers are unavailable, we retain the entire evaluation set. To ensure a fair evaluation, this filtering process is applied to all datasets and for all models, including our baselines. Applying this filter validates that datasets sourced from static knowledge bases like Wikipedia are particularly susceptible to this issue. For instance, we found that on HotpotQA and 2WikiMultiHopQA, the LLM already knew the answers to over 50% of

DATASET	BEFORE FILTERING	AFTER FILTERING	SAMPLE QUESTIONS
HotpotQA	7,405	2,291	What is the birthplace of the Senator who represents the first of 62 districts in the State Senate?
2WikiMultiHopQA	12,576	5,672	Which film came out first, Kauravar or Beyond The Sunset?
MuSiQue			
<i>Ans</i>	2,417	1,489	When did the people who captured Malakoff come to the region where Philipsburg is located?
<i>Full</i>	4,834	3,022	
CouldAsk	4,332	4,299	
<i>BBC</i>	278	264	What role did the Lord Chancellor hint might undergo further reforms in the next Labour manifesto?
<i>Yelp</i>	165	158	How is lost and found managed at the Charlotte Convention Center?
<i>Reddit</i>	313	301	Why was the student late to college yesterday?
<i>Squad v2</i>	1,000	1,000	How are the rates of social goods in countries with lower inequality?
<i>BanditQA</i>	2,070	2,070	How long are the Gospel of Matthew writings in the New Testament?
<i>QA2</i>	506	506	Why did the uk ban guns?
FaithEval	2,492	2,492	What is another word for the Earth’s upper mantle?

Table 2: Dataset statistics before and after filtering.

the questions. This rate of knowledge leakage was significantly lower in more recent or specialized answerability benchmarks.

## 4.2 Baseline Methods

Information sufficiency prediction is a nascent problem without much prior effort; we compare our framework against two recent works that use LLM prompting to assess sufficiency via two distinct strategies:

**(1) Sufficient Context Auto-rater** We use the LLM-based auto-rater from (Joren et al., 2025) by prompting GPT-4o in a one-shot setting to classify whether a given context is sufficient.

**(2) MIGRES** We adapt the framework from (Wang et al., 2024) by reformulating sufficiency detection as a question-answering task. The model is instructed to answer the question based on the provided information. If it cannot, it must respond with “unanswerable” and provide a summary of the missing information. We use the unanswerable classification as the sufficiency prediction, representing a strategy where sufficiency is inferred from the model’s ability to perform the end task.

## 4.3 Implementation Details

For the hypothesis generation (**Identify**) step, we use the smaller, open-source Llama-3.1-8B-Instruct (Meta, 2024) model. We generate five hypotheses for each input by setting the sampling temperature to 0.5. To establish semantic consensus, we use all-MiniLM-L6-v2 SentenceTransformer (Reimers and Gurevych, 2019) to encode the hypotheses into embeddings. A minimum similarity threshold of 0.3 is used to determine if a clear consensus exists.

For attribution verification (**Verify**), we use GPT-4o (OpenAI, 2024), which determines if claimed missing information is present, accepting synonyms, paraphrasing, and reasonable inferences as valid evidence. Notably, our modular pipeline generalizes across model families: the identification stage benefits from sampling diversity (achievable even with smaller models), while the verification task is straightforward enough that both open-source and closed-source LLMs perform comparably (see Section 5.5). The self-consistency checks run in parallel, and the consensus-finding step is a lightweight similarity computation, making the pipeline practical for real-world RAG applications where skipping sufficiency checks can lead to hallucinations or costly retrieval loops.

## 4.4 Evaluation Metrics

Since we have gold labels indicating whether the context is sufficient to answer a given question, we report accuracy to validate our approach. Furthermore for the FaithEval benchmark, since we have access to the ground-truth justification of why the context is insufficient, we use this to see how well our method’s missing information detection aligns with this justification. We use LLM (GPT-4o) as an auto-evaluator to rate both the explanations on a scale of 1-5, with 1 being not aligned (completely different and unrelated) and 5 being perfectly aligned (both state the same reason).

## 5 Results and Discussion

Our experiments are designed to answer five key research questions. **(RQ1)** How does our framework perform on complex inferential tasks compared to

Dataset	Accuracy (%)		
	Autorater	MIGRES	Ours
HotpotQA	83.99	84.47	<b>85.24</b>
2WikiMultihopQA	87.75	85.93	<b>89.31</b>
MuSiQue-Ans	64.85	66.96	<b>71.86</b>
MuSiQue-Full	59.49	68.65	<b>72.13</b>
<i>CouldAsk Benchmark</i>			
BBC	63.89	66.91	<b>78.78</b>
Yelp	61.81	56.36	<b>68.48</b>
Reddit	59.49	58.15	<b>69.65</b>
QA2	67.93	77.03	<b>78.46</b>
BanditQA	80.65	81.82	<b>83.33</b>
SQuAD v2	<b>82.87</b>	82.10	69.50
FaithEval	<b>78.08</b>	69.04	57.61

Table 3: Comparison of accuracy across QA datasets using different baselines. Green highlights best accuracies; second best are in red.

the baselines? **(RQ2)** Can the framework identify multiple and diverse types of missing information? **(RQ3)** How does the framework adapt to different sufficiency criteria, from pragmatic inference to strict literal matching? **(RQ4)** What is the role of the verification step? **(RQ5)** What is the impact of model choice and size on each pipeline stage?

### 5.1 RQ1: Performance Across Reasoning Complexity and Inferential Questions

Table 3 presents our main results. The advantage of our Identify-then-Verify framework grows in proportion to task reasoning complexity. On datasets like HotpotQA and 2WikiMultiHopQA, which are largely extractive in nature, our framework achieves competitive performance with marginal gains over baselines. This aligns with our hypothesis: when sufficiency reduces to verifying specific facts, simple prompting strategies suffice. However, on more complex datasets, the benefits are substantial. For MuSiQue, which features longer reasoning chains, our method yields a notable accuracy gain of up to 5 – 7 points. On the CouldAsk benchmark, our method shows superior performance, particularly on BBC, Yelp, and Reddit datasets (10+ point gains). These feature open-ended, inferential questions where the notion of *missing* information is ambiguous. The self-consistency step allows exploration of multiple interpretations, while the verifier consolidates these hypotheses into robust sufficiency judgement.

However, on SQuAD v2 and FaithEval, a simple LLM auto-rater for sufficiency works much better, while our method performs substantially poorly. Upon closer analysis, we note that this points to

a much bigger issue in sufficiency prediction, the definition of sufficiency and its subjective nature. Our framework’s default setting, which allows for reasonable inferences, is thus penalized by benchmarks requiring strict, literal fact-checking. We address this issue in detail in Section 5.3.

We also assess statistical significance to validate our gains using McNemar’s test, which is designed for comparing classifiers on the same dataset by analyzing their disagreements. Our improvements are statistically significant with  $p < 0.008$  on HotpotQA and 2WikiMultihopQA, and  $p < 0.0002$  on more complex datasets like MuSiQue and BBC. Per-class precision, recall, and F1 scores are provided in Appendix A.1.

### 5.2 RQ2: Identifying Multiple and Diverse Information Gaps

A key advantage of our generative formulation is the ability to identify not just whether information is missing, but *what* is missing. Unlike binary classification, our approach produces explicit, comprehensive descriptions of information gaps. Critically, a single hypothesis can identify multiple distinct missing pieces together (Table 4), and when multiple runs consistently generate similar multi-gap hypotheses, the semantic consensus captures this comprehensive assessment.

<b>Question</b>	Will Hilton scale NoMad internationally as quickly as it scaled Canopy and Tempo?
<b>Context</b>	Hilton acquired majority interest in Sydell Group to expand NoMad brand globally... [acquisition details, NoMad London flagship mentioned]
<b>MIGRES</b>	<b>Fails to flag insufficiency.</b> Proceeds to answer “Yes” confidently, citing “Hilton’s network of 6,100 hotels” and naming executives not present in the context. <i>The model generates a plausible-sounding response using parametric knowledge despite lacking critical information for proper inference.</i>
<b>Our Method</b>	<b>Correctly identifies insufficiency.</b> Consensus hypothesis identifies multiple missing pieces needed for inference: (1) specific strategies for scaling NoMad, (2) information about potential obstacles, (3) comparison to Canopy/Tempo expansion timelines. High semantic similarity across runs confirms this comprehensive assessment.

Table 4: Answering-based methods fail to detect insufficiency and generate confident responses using parametric knowledge, while our method correctly identifies multiple missing information gaps.

### 5.3 RQ3: Adapting to the Spectrum of Sufficiency

We observe that on SQuAD v2 and FaithEval, our default pragmatic configuration is suboptimal (Table 3). Our error analysis reveals that these benchmarks test strict contextual faithfulness, penalizing any form of inference (Table 5). For instance, when a question asks for the timing of a CEO’s resignation and the context states “The CEO announced he would step down next year,” our verifier deems it *Sufficient*, while the benchmark requires a precise date. This highlights that information sufficiency is not monolithic but exists on a spectrum, from pragmatic inference to literal fact-checking. Our framework accommodates this spectrum by treating verification strictness as a tunable parameter. For applications requiring faithfulness to exact wording, we simply reconfigure the verification prompt to reject inferences and require explicit matches. The *Identify* step remains unchanged; only the verification criteria shift.

<b>Question:</b>	What is something that teachers are at a high risk for?
<b>Key Context:</b>	Teachers face several occupational hazards in their line of work, including <i>occupational stress</i> , which can negatively impact teachers’ mental and physical health, productivity, and students’ performance. Stress can be caused by organizational change, relationships with students, fellow teachers, and administrative personnel, working environment, expectations to substitute, long hours with a heavy workload, and inspections.
<b>Ground Truth:</b>	<b>Insufficient</b>
<i>GT Justification:</i>	The sentence mentioning that teachers are at high risk for occupational burnout was removed. Therefore, the modified context no longer supports the original answer.
<b>Our prediction:</b>	<b>Sufficient</b>
<i>Our Reasoning:</i>	The context explicitly states that “Teachers face several occupational hazards...”, and directly names occupational stress as one of them.
<b>Question:</b>	Who owns the rail lines in Victoria?
<b>Key Context:</b>	Rail transport in Victoria is provided by <i>several private and public railway operators</i> . Major operators include: Metro Trains Melbourne, which runs an extensive, electrified, passenger system throughout Melbourne and suburbs; V/Line, which operates a concentrated service to major regional centres and long-distance services; Pacific National and CFCL Australia, which operate freight services.
<b>Ground Truth:</b>	<b>Insufficient</b>
<i>GT Justification:</i>	The sentence mentioning that the rail lines are ‘government-owned’ has been removed. Without this information, the ownership of the rail lines in Victoria becomes unclear from the context.
<b>Our prediction:</b>	<b>Sufficient</b>
<i>Our Reasoning:</i>	The passage mentions that rail transport in Victoria is provided by several private and public railway operators, which implies ownership or operational control by these entities. While it does not explicitly state ownership, it is reasonable to infer that these operators are responsible for the rail lines they operate.

Table 5: Two examples from FaithEval, where context is deemed **Insufficient** because it lacks the precise, literal answer, even though a reasonable inference is possible.

With this reconfiguration, our framework’s accuracy improves by over 15% on FaithEval (Table 6). Furthermore, the semantic alignment between our explanations and ground-truth justifications is 4.02 out of 5, compared to 3.89 for both baselines, confirming the model identifies gaps for the right reasons. This shows that a single unified method can adapt to different task requirements, from pragmatic QA systems to strict faithfulness benchmarks, with some tuning.

Dataset	Accuracy (%)		
	Autorater	MIGRES	Ours
SQuAD v2	82.87	82.10	<b>84.06</b>
FaithEval	78.08	69.04	<b>90.24</b>

Table 6: Insufficiency accuracy with leniency ablation.

### 5.4 RQ4: The Role of Verification

To measure the contribution of the verification step, we compare our full Identify-then-Verify framework against an Identify-Only baseline that makes predictions based solely on semantic consensus. Figure 4 shows the Relative Error Reduction (RER) achieved by adding verification. We use RER rather than absolute accuracy gain as it measures the percentage of errors from the Identify-Only step that were successfully corrected. The verification step yields significant RER across most datasets, with gains as high as 37%. This confirms that the verification step is a critical self-correction mechanism, effectively catching hallucinations or misinterpretations from the initial identification stage.

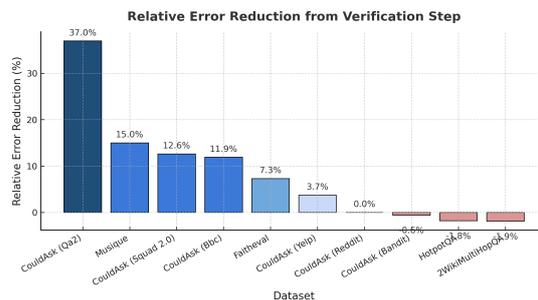


Figure 4: Relative Error Reduction (RER) from the verification step across various datasets. The bars show the percentage of errors from the initial consensus that were corrected.

### 5.5 RQ5: Model Ablation Study

To demonstrate the generalizability of our framework across model families and scales, we ablate

Dataset	Baselines		Ours (+ GPT-4o Verify)	
	Auto.	MIG.	Llama 8B	GPT-4o
HotpotQA	83.99	84.47	85.24	<b>87.60</b>
2WikiMHQA	87.75	85.93	89.31	<b>92.60</b>
MuSiQue	64.85	66.96	71.86	<b>73.00</b>
BBC	63.89	66.91	<b>78.78</b>	77.34
Yelp	61.81	56.36	68.48	<b>71.52</b>
Reddit	59.49	58.15	69.65	<b>70.60</b>
BanditQA	80.65	81.82	<b>83.33</b>	82.95
FaithEval	78.08	69.04	<b>90.24</b>	89.62

(a) Effect of varying the Identify model.

Dataset	Baselines		Ours (Llama 8B Identify +)		
	Auto.	MIG.	GPT-4o	Claude 4	Llama 8B
HotpotQA	83.99	84.47	<b>85.24</b>	84.52	84.64
2WikiMHQA	87.75	85.93	<b>89.31</b>	87.87	88.00
MuSiQue	64.85	66.96	<b>71.86</b>	69.91	59.20
BBC	63.89	66.91	78.78	<b>79.50</b>	71.14
Yelp	61.81	56.36	<b>68.48</b>	64.85	63.52
Reddit	59.49	58.15	<b>69.65</b>	65.00	63.67
BanditQA	80.65	81.82	<b>83.33</b>	78.89	79.03
FaithEval	78.08	69.04	<b>90.24</b>	88.20	86.00

(b) Effect of varying the Verification model.

Table 7: Model ablation study across pipeline stages. (a) Even a smaller 8B model for identification nearly matches GPT-4o while significantly outperforming baselines. (b) The verification task is straightforward enough that both open-source and closed-source LLMs perform comparably.

both pipeline stages on a representative subset of our evaluation datasets.

**Effect of Identification Model.** Table 7a reveals that even the smaller Llama-3.1-8B-Instruct model generates high-quality hypotheses for the identification stage, with performance nearly matching GPT-4o while significantly outperforming baselines. This validates our design choice of using a cost-effective smaller model for the first stage.

**Effect of Verification Model.** Table 7b shows that while fixing identification to Llama-3.1-8B-Instruct, various LLMs (GPT-4o, Claude 4 Sonnet, and Llama-3.1-8B-Instruct) perform comparably on the verification task, all consistently exceeding baselines. This confirms that verification is a straightforward enough task that for both open-source and closed-source models.

Overall, these results validate the effectiveness of our task formulation across diverse model types, allowing practitioners to choose models based on cost or latency requirements.

## 6 Conclusion

We address robust assessment of information sufficiency for question answering with LLMs. While existing approaches rely on direct binary classification or answering-based detection, both prove brittle for complex inferential questions. We propose an Identify-then-Verify framework that reformulates sufficiency as a generative task: explicitly identifying what information is missing. Our method generates multiple hypotheses via self-consistency, establishes semantic consensus, and performs verification to guard against hallucination. This yields three key advantages: (1) improved accuracy on complex inference questions; (2) comprehensive identification of multiple information gaps, providing actionable guidance for retrieval;

and (3) adaptability to different sufficiency criteria, from pragmatic inference to strict literal matching, through simple prompt reconfiguration. A central insight is that information sufficiency varies with application requirements and user needs. By treating verification strictness as a tunable parameter, our framework accommodates diverse requirements. This flexibility, combined with explicit gap identification, makes our approach more reliable and interpretable, paving the way for trustworthy QA systems that know when to abstain and why.

## Limitations

Our self-consistency step introduces a trade-off between accuracy and computational cost. While we mitigate this by using a smaller model for hypothesis generation, future research could explore more efficient sampling strategies. Further, the adaptability of our framework is currently guided by a manually set strictness definition for sufficiency. A promising direction for future work would be to develop methods that automatically determine the optimal verification setting for a given task, making the framework even more autonomous. Finally, while existing benchmarks serve as valuable initial testbeds for sufficiency assessment, they often do not fully capture the complexity of real-world information-seeking scenarios. Many real-world questions, such as the example in Table 4, are inherently open-ended and inferential, requiring models to navigate incomplete contexts with multiple potential information gaps. Developing benchmarks that better reflect these characteristics would benefit the broader research community.

## Ethics Statement

There are no ethical concerns to the best of our knowledge.

## References

- Lu Dai, Yijie Xu, Jinhui Ye, Hao Liu, and Hui Xiong. 2025. [Seper: Measure retrieval utility through the lens of semantic perplexity reduction](#).
- Yang Deng, Yong Zhao, Moxin Li, See-Kiong Ng, and Tat-Seng Chua. 2024. [Don't just say "i don't know"! self-aligning large language models for responding to unknown questions with explanations](#).
- Ekaterina Fadeeva, Aleksandr Rubashevskii, Roman Vashurin, Shehzaad Dhuliawala, Artem Shelmanov, Timothy Baldwin, Preslav Nakov, Mrinmaya Sachan, and Maxim Panov. 2025. [Faithfulness-aware uncertainty quantification for fact-checking the output of retrieval augmented generation](#).
- Chenrui Fan, Ming Li, Lichao Sun, and Tianyi Zhou. 2025. [Missing premise exacerbates overthinking: Are reasoning models losing critical thinking skill?](#)
- Kevin Fischer, Darren Fürst, Sebastian Steindl, Jakob Lindner, and Ulrich Schäfer. 2024. [Question: How do large language models perform on the question answering tasks? answer:.](#)
- Harvey Yiyun Fu, Aryan Shrivastava, Jared Moore, Peter West, Chenhao Tan, and Ari Holtzman. 2025. [Absencebench: Language models can't tell what's missing](#).
- Ge Gao, Eunsol Choi, and Yoav Artzi. 2022. [Simulating bandit learning from user feedback for extractive question answering](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5167–5179, Dublin, Ireland. Association for Computational Linguistics.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. [Constructing a multi-hop qa dataset for comprehensive evaluation of reasoning steps](#).
- Akriti Jain and Aparna Garimella. 2025. [Modeling contextual passage utility for multihop question answering](#). In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 464–471, Mumbai, India. The Asian Federation of Natural Language Processing and The Association for Computational Linguistics.
- Ziqi Ji, Nayeon Lee, Rita Frieske, et al. 2023. [Survey of hallucination in natural language generation](#). *ACM Computing Surveys*.
- Hailey Joren, Jianyi Zhang, Chun-Sung Ferng, Da-Cheng Juan, Ankur Taly, and Cyrus Rashtchian. 2025. [Sufficient context: A new lens on retrieval augmented generation systems](#).
- Hazel Kim, Adel Bibi, Philip Torr, and Yarin Gal. 2024. [Detecting llm hallucination through layer-wise information deficiency: Analysis of unanswerable questions and ambiguous prompts](#).
- Najoung Kim, Phu Mon Htut, Samuel R. Bowman, and Jackson Petty. 2023. [\(qa\)<sup>2</sup>: Question answering with questionable assumptions](#).
- Polina Kirichenko, Mark Ibrahim, Kamalika Chaudhuri, and Samuel J. Bell. 2025. [Abstentionbench: Reasoning llms fail on unanswerable questions](#).
- Kalpesh Krishna and et al. 2024. [Hallucination in retrieval-augmented generation: Benchmarking and mitigation](#). *arXiv preprint arXiv:2403.00000*.
- Ivano Lauriola, Stefano Campese, and Alessandro Moschitti. 2025. [Analyzing and improving coherence of large language models in question answering](#). In *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)*, pages 11740–11755, Albuquerque, New Mexico. Association for Computational Linguistics.
- Zichong Li, Xinyu Feng, Yuheng Cai, Zixuan Zhang, Tianyi Liu, Chen Liang, Weizhu Chen, Haoyu Wang, and Tuo Zhao. 2025. [Llms can generate a better answer by aggregating their own responses](#).
- Chuangtao Ma, Yongrui Chen, Tianxing Wu, Arijit Khan, and Haofen Wang. 2025. [Large language models meet knowledge graphs for question answering: Synthesis and opportunities](#).
- Nishanth Madhusudhan, Sathwik Tejaswi Madhusudhan, Vikas Yadav, and Masoud Hashemi. 2024. [Do llms know when to not answer? investigating abstention abilities of large language models](#).
- Meta. 2024. [The llama 3 herd of models](#). *arXiv preprint arXiv:2407.21783*.
- Yifei Ming, Senthil Purushwalkam, Shrey Pandit, Zixuan Ke, Xuan-Phi Nguyen, Caiming Xiong, and Shafiq Joty. 2025. [Faitheval: Can your language model stay faithful to context, even if "the moon is made of marshmallows"](#).
- OpenAI. 2024. [Gpt-4 technical report](#).
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don't know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

- Asir Saadat, Tasmia Binte Sogir, Md Taukir Azam Chowdhury, and Syem Aziz. 2024. [When not to answer: Evaluating prompts on gpt models for effective abstention in unanswerable math word problems.](#)
- Zhengyan Shi, Giuseppe Castellucci, Simone Filice, Saar Kuzi, Elad Kravi, Eugene Agichtein, Oleg Rokhlenko, and Shervin Malmasi. 2025. [Ambiguity detection and uncertainty calibration for question answering with large language models.](#) In *Proceedings of the 5th Workshop on Trustworthy NLP (TrustNLP 2025)*, pages 41–55, Albuquerque, New Mexico. Association for Computational Linguistics.
- Christian Tomani, Kamalika Chaudhuri, Ivan Evtimov, Daniel Cremers, and Mark Ibrahim. 2024. [Uncertainty-based abstention in llms improves safety and reduces hallucinations.](#)
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. [Musique: Multihop questions via single-hop question composition.](#)
- Keheng Wang, Feiyu Duan, Peiguang Li, Sirui Wang, and Xunliang Cai. 2024. [Llms know what they need: Leveraging a missing information guided framework to empower retrieval-augmented generation.](#)
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models.](#)
- Roy Xie, Junlin Wang, Paul Rosu, Chunyuan Deng, Bolun Sun, Zihao Lin, and Bhuwan Dhingra. 2025. [Knowing when to stop: Dynamic context cutoff for large language models.](#)
- Yilong Xu, Jinhua Gao, Xiaoming Yu, Yuanhai Xue, Baolong Bi, Huawei Shen, and Xueqi Cheng. 2025. [Training a utility-based retriever through shared context attribution for retrieval-augmented language models.](#)
- Yongjin Yang, Haneul Yoo, and Hwaran Lee. 2025. [Maqa: Evaluating uncertainty quantification in llms regarding data uncertainty.](#)
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [Hotpotqa: A dataset for diverse, explainable multi-hop question answering.](#)
- Murong Yue. 2025. [A survey of large language model agents for question answering.](#)
- Hengran Zhang, Keping Bi, Jiafeng Guo, Jiaming Zhang, Shuaiqiang Wang, Dawei Yin, and Xueqi Cheng. 2025. [Distilling a small utility-based passage selector to enhance retrieval-augmented generation.](#)
- Wenting Zhao, Ge Gao, Claire Cardie, and Alexander M. Rush. 2024. [I could've asked that: Reformulating unanswerable questions.](#)

## A Appendix

### A.1 Detailed Performance Metrics

Table 8 presents per-class precision (P), recall (R), and F1 scores for all methods.

Dataset	Class	Ours			MIGRES			Autorater		
		P	R	F1	P	R	F1	P	R	F1
HotpotQA	Insufficient	87.7	82.0	84.7	81.8	88.7	85.1	78.9	90.5	84.3
	Sufficient	83.1	88.5	85.7	87.7	80.2	83.8	88.8	75.8	81.8
MuSiQue-Ans	Insufficient	67.8	82.0	74.2	61.2	92.6	73.7	58.9	97.8	73.5
	Sufficient	77.2	61.1	68.2	84.8	41.3	55.6	93.7	31.9	47.6
2WikiMHQA	Insufficient	83.7	98.1	90.3	80.4	95.1	87.1	89.1	86.1	87.5
	Sufficient	97.7	80.9	88.5	94.0	76.8	84.5	86.5	89.4	87.9
BBC	Insufficient	50.0	47.5	48.7	33.3	55.9	41.8	33.3	69.5	45.1
	Sufficient	86.0	87.2	86.6	85.5	69.9	76.9	88.3	62.4	73.1
Yelp	Insufficient	49.1	51.0	50.0	39.2	74.5	51.4	43.8	82.4	57.1
	Sufficient	77.7	76.3	77.0	80.9	48.2	60.4	87.0	52.6	65.6
FaithEval	Insufficient	100.0	90.2	94.9	100.0	69.0	81.7	100.0	78.1	87.7

Table 8: Per-class precision, recall, and F1 scores. Our method generally achieves higher F1 scores, particularly on the Sufficient class where baselines often over-predict insufficiency.