HOW SHOULD WE EVALUATE LLM REASONING QUALITY FOR FACT VERIFICATION?

Anonymous authors

000

001

002 003 004

010 011

012

013

014

016

018

019

021

023

024

025

026

027

028

029

031

034

037

039

040 041

042

043

044

046 047

048

051

052

Paper under double-blind review

ABSTRACT

The reasoning traces generated by Large Language Models (LLMs) are increasingly used to improve final predictions, enable reinforcement learning based on reasoning trace correctness, and justify model outputs to users. Their recognized utility spurred a line of works on evaluating LLM reasoning quality. However, such current reasoning evaluation methods are typically generic and do not shed light on the different reasoning types that may be required for various complex tasks. In this paper, we investigate reasoning quality for the prominent task of Fact Verification, where a model should determine whether a given claim is entailed by a reference source text, a fundamental process known as Natural Language Inference (NLI). Specifically, we propose a novel evaluation framework that considers the prominent types of inference steps involved in NLI reasoning: hypothesis decomposition into individual facts, followed by source attribution and entailment decision for each fact, and finally aggregation of fact level decisions into the final entailment classification. Our protocol introduces fine-grained metrics to assess both the existence (whether a step was performed) and the quality (how well it was performed) for each inference type. Following this framework, we first conduct a meticulous manual evaluation of six prominent LLMs, and then scale the evaluation using LLM-as-a-Judge. Our analysis reveals several insights, including: (1) a significant positive correlation exists between the quality of the reasoning trace and the correctness of the final prediction; (2) models often omit necessary reasoning steps, leading to incomplete justifications; and (3) guiding the LLM towards a systematic reasoning trace based on our framework often improves the quality of both the reasoning trace and the overall entailment classification, specifically for "non-reasoning" models. Overall, our work provides a more diagnostic and nuanced approach to understanding and evaluating LLM reasoning trace, demonstrated specifically for NLI reasoning in fact verification, proposing insights for future improvements in reasoning quality and its downstream usage.

1 Introduction

Explicit reasoning traces have important benefits like improving final answer prediction (Wei et al., 2022; Sprague et al., 2025); guiding search algorithms over potential reasoning paths to improve model predictions (Hao et al., 2024; Sun et al., 2024); providing granular feedback for fine-tuning models via reinforcement learning, where the reward is based on the quality of intermediate steps (Lai et al., 2024; Lu et al., 2024; Sun et al., 2024); and might serve as a verification and justification of the final predictions in human-AI collaboration scenarios (Barez et al., 2025).

The growing importance of these reasoning traces has spurred research into methods for evaluating their quality. Existing approaches typically fall into two categories. Some provide a single, holistic score for an entire reasoning trace (Saparov & He, 2023; He et al., 2024), while others adopt a more granular, step-by-step analysis, evaluating generic properties of inference steps, such as ground-edness or faithfulness (Prasad et al., 2023; Li et al., 2025). However, a key limitation of these step-by-step methods is that they often apply a uniform set of criteria to every step, disregarding the distinct functional role that each type of step plays within a larger, structured reasoning process (like those elaborated below for NLI reasoning).

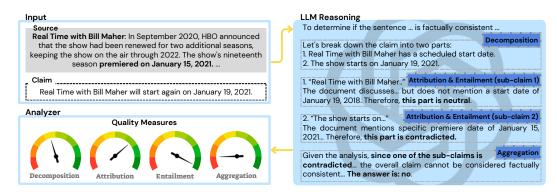


Figure 1: Overview of our evaluation framework. A model produces a reasoning trace before its final prediction on the verification of a given claim. Our framework then evaluates the quality of this trace across four distinct components: Decomposition, Attribution, Entailment, and Aggregation.

This limitation is particularly salient for complex, multi-stage tasks. In this work, we focus on entailment classification, specifically in the prominent context of fact verification, as often applied either to LLM- or human-generated claims (Tian et al., 2020; Thorat et al., 2025; Ádám Kovács & Recski, 2025; Paudel et al., 2025). The underlying reasoning required in this setting is Natural Language Inference (NLI, aka as textual entailment), where a model should determine if a reference text (*premise*) entails a given claim (*hypothesis*) (Dagan et al., 2005; Bowman et al., 2015). A robust reasoning process for NLI typically involves several distinct types of reasoning steps: decomposing the claim into sub-claims, attributing each sub-claim to supporting or refuting evidence in the source, classifying the entailment status of each sub-claim, and finally, aggregating these individual judgments into a final verdict for the entire original claim (see Fig. 1).

Building on this structured view of NLI reasoning, we introduce a novel evaluation protocol designed to assess a model's proficiency at each distinct type of reasoning steps. Our scheme evaluates each component – from decomposition to aggregation – along two axes: its *existence* (i.e., whether the model included the step in the reasoning trace) and its *quality* (e.g., the correctness of the step). This protocol is designed for both human evaluation as well as for scalable automated evaluation using an LLM-as-a-Judge.

Our experiments reveal that models often fail to produce a complete and valid reasoning trace, and often omit some necessary steps. Often, guiding the LLM to generate a complete reasoning trace, consisting of all required inference stages, successfully encourages this behavior, and in some cases improves the quality of both reasoning and final task accuracy. Furthermore, our analysis confirms a substantial correlation between the quality of individual reasoning steps and the final prediction. Specifically, the correctness of the Attribution and Entailment components for each sub-claim are strong predictors of a correct final answer, increasing confidence in the model's conclusion when these steps are performed correctly.

Overall, our contributions include: (1) a novel evaluation methodology for NLI reasoning, that decomposes the complex, multi-stage task of fact verification into distinct, functionally-motivated components; (2) a set of fine-grained metrics for evaluating both the existence and the quality of each reasoning component, enabling a more precise diagnosis of model performance; (3) an extensive empirical analysis of six prominent LLMs, with both unguided and guided prompting, providing a detailed comparison of their capabilities at each stage of the NLI reasoning process.

2 BACKGROUND

Fact Verification. Fact verification was first formally defined as a computational task by Vlachos & Riedel (2014), where the goal is to assess the faithfulness of a given claim. This process typically involves two main stages: retrieving a relevant evidence document and then verifying the claim against that evidence. The focus of our project is on the verification step, assuming an evidence document has already been retrieved. While the main body of fact verification research has focused on textual evidence (Wang, 2017; Thorne et al., 2018; Kamoi et al., 2023; Schuster et al., 2021),

a few studies have also built datasets with other evidence types, such as tables and images (Chen et al., 2020; Yao et al., 2023). Many approaches have been proposed for the verification step itself. A dominant paradigm is to frame the problem as an entailment decision (Dagan et al., 2005), where the system must determine if the evidence supports or refutes the claim. A few methods relied on fine-tuned models (Yang et al., 2021; Chen et al., 2022; Tang et al., 2024), while more recent approaches have leveraged LLMs (Zeng & Gao, 2023; Li et al., 2024; Parvez, 2025). We focus on this latter approach, specifically on cases of prompting LLMs to produce explicit reasoning tokens Lei et al. (2023); Wadhwa et al. (2024); Wan et al. (2025).

Reasoning Steps Usefulness. The role and usefulness of intermediate reasoning text, aka Chain-of-Thought tokens, is a subject of ongoing debate in the community. A significant body of work argues that chain-of-thought (CoT) tokens are not faithful explanations of a model's decision-making process; that is, they do not necessarily reflect the internal computations that produce the final answer (Kambhampati et al., 2025; Barez et al., 2025). These studies show that reasoning traces may be partially incorrect, omit critical information, or present a plausible but fabricated justification for an incorrect prediction (Turpin et al., 2023; Stechly et al., 2025; Bhambri et al., 2025).

At the same time, there is compelling counterevidence that higher-quality reasoning traces demonstrably improve final task performance (Liao et al., 2025; Gandhi et al., 2025). Their practical utility has been shown in several applications, such as guiding search over multiple reasoning paths or providing reward signals for fine-tuning models via reinforcement learning (Hao et al., 2024; Sun et al., 2024; Lai et al., 2024; Lu et al., 2024). This has led to the perspective that CoT should be treated as a valuable communication tool, provided its quality is systematically evaluated (Barez et al., 2025). Our work is motivated by this view: rather than demanding strict faithfulness to the model's internal process, we propose a protocol to rigorously evaluate the reasoning trace's quality for the downstream task of justifying an NLI decision.

Reasoning Steps Evaluation. Prior evaluations of reasoning traces proceed along three main lines. (i) *A unified score* methods score the entire reasoning process as a whole (Saparov & He, 2023; Han et al., 2024; He et al., 2024). (ii) *Reference-based* approaches compare each generated step to some gold steps (Hao et al., 2024; Li et al., 2025). (iii) *Reference-free* approaches propose generic, step-level criteria – including groundedness, semantic consistency, logical validity, fluency, minimality, and efficiency (Prasad et al., 2023; Golovneva et al., 2023; Saparov & He, 2023; Zhou et al., 2025; Qiu et al., 2025; Chen et al., 2025; Li et al., 2025).

While flexible, the primary limitation of most of the existing reference-free metrics is their generic nature. They either treat each step in isolation or, at best, check for local consistency with the preceding steps. Crucially, they neglect the fact that steps within a complex reasoning process serve distinct functions, yet they apply a uniform set of evaluation criteria to all of them. This approach fails to provide a nuanced understanding of a model's capabilities at different stages of a task. In contrast, we argue for a type-aware evaluation. We propose to distinguish between step types and evaluate each based on its specific function. This methodology enables a more fine-grained analysis of a model's ability to execute the different kinds of reasoning steps necessary to produce a final, aggregated answer.

3 A PROTOTYPICAL NLI REASONING SCHEME

A prerequisite for assessing current models' entailment reasoning capabilities is understanding the expected structure of a complete reasoning process. A prominent available source for this purpose can be found in the annotation guidelines used in previous works that constructed NLI and fact verification datasets. These guidelines often instruct annotators to follow a systematic workflow before deciding on the entailment judgment of a claim. Additionally, we are inspired by established system architectures for entailment classification models, where often different components were responsible for different types of reasoning steps.

A common first step, particularly relevant in real-world scenarios where the given claims are naturally occurring fairly long sentences, is to *decompose* the claim into smaller sub-claims, making it easier to verify each sub-claim separately (Min et al., 2023; Kamoi et al., 2023; Mishra et al., 2024; Mitra et al., 2025). This decomposition necessitates two subsequent steps: determining the

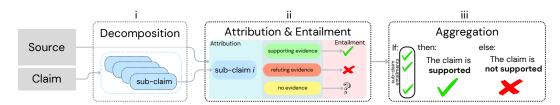


Figure 2: Overview of a comprehensive NLI reasoning process. (i) Decomposition: the original claim is split into individual sub-claims. (ii) Attribution & Entailment: each sub-claim is checked against the source for supporting evidence, refuting evidence, or no evidence. (iii) Aggregation: if all sub-claims are supported, the claim is accepted; otherwise, it is rejected.

entailment status for each sub-claim (i.e., whether it is supported or not by the given source, optionally distinguishing contradiction vs neutral cases), and then an aggregation of these individual judgments into a final decision. The aggregation logic dictates that the entire claim is supported only if all sub-claims are entailed; otherwise, it is considered not supported. As a special case, if no decomposition occurs, the entailment decision is simply made for the claim as a whole, where no aggregation is needed.

The reasoning process must also include a meaningful step of *attribution* (i.e. evidence detection), where the model searches the source for evidence that either supports or refutes a sub-claim (Camburu et al., 2018; Niu et al., 2024; Wang & Atanasova, 2025). If such evidence is found the sub-claim entailment status is classified accordingly (entailed/supported or contradicted/refuted); otherwise the entailment status is neutral/unknown.

Finally, the correctness of a sub-claim is not always directly derivable from the evidence but may require some *inference* (Camburu et al., 2018; Bhagavatula et al., 2020; Niu et al., 2024; Havaldar et al., 2025). For example, given the sub-claim "The adult interrupted Donald Trump's speech" and the evidence "A 50-year-old man interrupted Donald Trump's speech", an inference that 'a 50-year-old man is an adult' leads to determining the entailment status. Thus, an inference step serves as a complementary bridge between attribution and the final entailment classification. The flow including the above reasoning stages is illustrated in Fig. 2.

4 NLI REASONING EVALUATION

The previous section outlined four key components of NLI reasoning: decomposition, attribution, entailment classification, and aggregation. Building on these, we propose several metrics that each assess a different aspect of the reasoning process. Taken together, these metrics provide a comprehensive analysis of NLI reasoning. The following metrics comprise the evaluation scheme for annotators, as detailed in Section 5.

We divide the evaluation metrics into two groups: *existence* metrics and *quality* metrics. The first group measures whether each component is present in the model's reasoning process, while the second group evaluates the quality and correctness of those components when they are executed.

4.1 Existence Metrics

The existence metrics are binary values that indicate, for each reasoning instance, whether a model produces each of the reasoning components: *decomposition* - whether the model decomposes the claim into smaller sub-claims; *attribution* - whether the model searches evidence for each sub-claim; *inference* - whether the model describe the inference required for an entailment classification for each sub-claim, *entailment* - whether the model determines the entailment status of each sub-claim and *aggregation* - whether the model aggregates the entailment decisions of all sub-claim (when there are a few sub-claims). The decomposition and aggregation metrics are binary values for each reasoning trace. In contrast, the other existence metrics are calculated as the proportion of sub-claims for which the component is present. Next, we describe the *quality* metrics for each component, which we only evaluate if the component exists in the reasoning trace of the respective instance.

4.2 DECOMPOSITION

Granularity. Inspired by (Wanner et al., 2024), we define the decomposition *granularity* for a reasoning instance as the number of distinct sub-claims $\mathcal{H} = \{h_1, h_2, \ldots, h_n\}$ generated at the decomposition step. If no decomposition occurs, \mathcal{H} contains a single element. The granularity score is then defined as: $G := |\mathcal{H}|$. This metric has no ground-truth value, but it can influence later steps. Low granularity leads to longer and more complex sub-claims, making attribution and entailment classification harder. High granularity increases the risk of unfaithful or incomplete decompositions. In Fig. 1 example, the granularity value is 2.

Soundness. As part of the decomposition step, we assess whether the model, in its reasoning steps, generates sub-claims that are semantically entailed by the claim. The *soundness* metric measures the proportion of generated sub-claims that are consistent with the claim. The soundness score (for a reasoning instance) is defined as: $S := \frac{1}{|\mathcal{H}|} \sum_{i=1}^{|\mathcal{H}|} \mathbb{1}_{\{h_i \text{ is sound}\}}$. Intuitively, a low soundness score suggests the model introduces extraneous or fabricated sub-claims during decomposition, risking incorrect entailment judgments. In Fig. 1, both sub-claims are sound, therefore the value of this metric is 1.

Completeness. For a complete view of the decomposition step, we evaluate whether the model refers all the semantic content of the original claim. The *completeness* metric checks if any part was omitted during decomposition. It is a binary value: 1 if all information is covered by the model's subclaims, and 0 if any is missing. The completeness score is then defined as: $C := \begin{cases} 1 & \text{if } \mathcal{H} \subseteq \bigcup_i h_i \\ 0 & \text{otherwise} \end{cases}$. Intuitively, this metric highlights cases where the model omits parts of the claim, potentially leading to incorrect predictions like falsely labeling it as *entailed*. In Fig. 1, there is no missing information, resulting in a value of 1.

4.3 ATTRIBUTION

Attribution Correctness. The second component in a comprehensive reasoning trace begins with an attribution for each sub-claim. This metric assesses whether the model correctly identifies supporting or contradicting evidence from the source, or indicates that no evidence exists, for each sub-claim. An attribution is considered correct if it justifies the sub-claim's entailment label. Formally, the metric is defined as: $A_{\rm tt} := \frac{1}{|\mathcal{H}|} \sum_{i=1}^{|\mathcal{H}|} \mathbbm{1}_{\{h_i \text{ is correctly attributed}\}}$. Intuitively, missing or incorrect attribution can cause sub-claim misclassification, leading to an incorrect overall entailment decision. For instance, in Fig. 1, the incorrect attribution for one of the two sub-claims results in a score of 1/2.

Furthermore, we distinguish between three types of attribution: (i) **extractive** – the model copy the exact evidence span(s) from the document; (ii) **paraphrase** – the model does not extract the exact span but paraphrases the relevant content from the document; and (iii) **abstract** – the model provides a higher-level explanation of the relevant information in the document. While the specific type of attribution is not necessarily critical for the correctness of the reasoning trace, it might affect the usefulness of the reasoning trace. For example, an extractive attribution can assist a human in verifying the final answer more easily than an abstractive attribution can. Therefore, this distinction serves an analytical purpose and should not be interpreted as a measure of quality.

4.4 Entailment

The second phase in the *Attribution & Entailment* step is to determine the entailment status of each sub-claim. If the model has produced an attribution, it may also perform an inference step, as described in Section 3.

Entailment Inference Correctness. This metric evaluates the correctness of an entailment inference, when it exists. $I_{\text{nfer}} := \frac{1}{|\mathcal{H}|} \sum_{i=1}^{|\mathcal{H}|} \mathbbm{1}_{\{h_i \text{ is correctly inferred}\}}$

¹If no sub-claims are generated, we treat the whole claim as a single sub-claim.

Entailment Classification Correctness. This metric evaluates whether the model correctly predicts the entailment label for each sub-claim by comparing the predicted label \hat{y}_i with the gold label y_i (provided by an oracle or human annotator). $E_{\text{ntail}} := \frac{1}{|\mathcal{H}|} \sum_{i=1}^{|\mathcal{H}|} \mathbb{1}_{\{\hat{y}_i = y_i\}}$ Intuitively, misclassifying even a single sub-claim can affect the overall claim prediction, making this step crucial for reliable performance. In Fig. 1, the model makes an inference and classification decision for both sub-claims, but since the inference and the classification of the first sub-claim are incorrect, the corresponding correctness metrics are both 1/2.

4.5 AGGREGATION

Aggregation Correctness. This metric is a binary score (0 or 1) that verifies if the model's final decision logically follows from its judgments on the sub-claims. The aggregation is considered correct if the final prediction is supported only when all sub-claims are supported, and not supported otherwise. This metric specifically captures inconsistencies between the reasoning trace and the final answer; for example, the correct aggregation in Fig. 1 earns a score of 1.

Overall Entailment Classification. The Overall Entailment Judgment is the final, bottom-line metric that measures the accuracy of the model's prediction against the gold-standard label, primarily used to analyze the correlation between reasoning quality and final task accuracy. A key property of our framework is its diagnostic completeness: an incorrect final prediction must originate from a flaw in at least one of the preceding reasoning components, ensuring the error is captured by our metrics. For example, the correct final answer in Fig. 1 results in a score of 1 for this metric.

5 EVALUATION & ANALYSIS

We begin by describing the experimental setup for our manual evaluation and then present the results of this meticulous evaluation over a representative sample, discussing key insights for both the *existence* and *quality* metrics. Following, we conduct a correlation analysis to investigate the relationships between the different reasoning components. Finally, we scale up our evaluation using a validated LLM-as-a-Judge to analyze the entire dataset, which allows us to assess a broader generalizability of our human evaluation findings.

5.1 EXPERIMENTAL SETUP

To analyze the reasoning traces produced by LLMs, we evaluated six prominent models: Llama-3-1B, Llama-3-8B, Llama-3-70B, Gemini-2.0-Flash, Gemini-2.5-Flash, and DeepSeek-Rl-32B-Distill. The latter two are reasoning models, which are specifically optimized for reasoning tasks. For our evaluation, we randomly selected 30 samples from the recent ClearFacts dataset, which combined 14 different fact-checking benchmarks (Seo et al., 2025). We tested two prompt variants for each model: (i) an *unguided CoT prompt*, which allows the model to generate its own reasoning structure to reach a final prediction, and (ii) a *guided CoT prompt*, where we explicitly instruct the model to follow the structured three-steps reasoning process described in Section 3. The full prompts are provided in Appendix F.

The generated outputs were evaluated by three annotators all having extensive experience in similar annotation tasks. Prior to the main annotation, they completed a dedicated training session on our specific evaluation protocol. In total, our manual evaluation comprises 360 annotated reasoning instances (30 samples \times 6 models \times 2 prompt variants).

A detailed inter-annotator agreement analysis, described in Appendix A, shows strong results on Gwet's score (Gwet, 2008) (score range between -1 and 1): 10 of our 13 metrics achieve 'Almost Perfect' agreement, while 3 are rated as 'Substantial' – according to the common interpretation of Landis & Koch (1977). This validates the high-quality data used in our main evaluation.

²For binary classification, the *neutral* and *contradicted* classes may be merged into a single *not supported* class.

Table 1: Manual evaluation results on the *quality metrics* across models. The results are the average scores across all the sampled instances. Results are split into (a) decomposition and attribution, and (b) entailment, aggregation, and overall correctness. Each cell shows unguided CoT (left) vs. guided CoT (right). The highest score in each column is **bold**. The two models below the horizontal line are reasoning models.

Model		Attribution		
	Granularity	Soundness	Completeness	
Llama-3-1B	1.56 / 5.04	0.41 / 0.46	0.38 / 0.61	0.55 / 0.71
Llama-3-8B	2.32 / 2.87	0.97 / 0.82	0.93 / 0.93	0.97 / 0.89
Llama-3-70B	2.54 / 2.73	0.97 / 0.89	0.88 / 0.90	0.93 / 0.86
Gemini-2.0-Flash	2.11 / 2.33	1.00 / 0.96	0.67 / 1.00	0.88 / 0.90
DeepSeek-R1-distill Gemini-2.5-Flash	2.15 / 2.93 2.20 / 3.03	0.93 / 0.86 1.00 / 0.93	0.95 / 0.86 0.88 / 0.88	0.96 / 0.90 1.00 / 0.96

Model	Inference	Entailment	Aggregation	Overall
Llama-3-1B	0.40 / 0.79	0.70 / 0.73	1.00 / 0.75	0.40 / 0.69
Llama-3-8B	0.82 / 0.81	0.90 / 0.83	0.91 / 1.00	0.73 / 0.75
Llama-3-70B	0.86 / 0.85	0.89 / 0.85	0.92 / 1.00	0.87 / 0.90
Gemini-2.0-Flash	0.80 / 0.99	0.75 / 0.99	0.89 / 1.00	0.77 / 0.90
DeepSeek-R1-distill Gemini-2.5-Flash	0.92 / 0.98 0.95 / 0.96	0.86 / 0.85 0.94 / 0.96	0.95 / 0.97 1.00 / 1.00	0.83 / 0.80 0.90 / 0.87

5.2 MANUAL EVALUATION RESULTS

We begin with a manual evaluation of a representative sample of model outputs. This evaluation allows us to establish a high-quality, gold-standard dataset for our core findings.

5.2.1 Existence Metrics

The average results of our manual evaluation for the existence metrics, across all the models, are presented in Fig. 3. The figure details the frequency with which each reasoning component was included in the reasoning trace, comparing the performance of unguided CoT (left of each cell) and guided CoT (right of each cell) for each model. The full results are presented in Appendix B, Table 3.

We observe, from both tables, several key trends. First, models frequently omit the various types of reasoning steps when unguided. This tendency is most pronounced for the decomposition step, which models often skip entirely unless explicitly prompted. Second, model scale appears to be a factor; larger models are more likely to spontaneously perform these steps even without guidance, though this

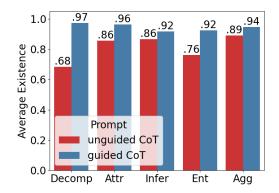


Figure 3: Manual evaluation average results on the *existence metrics* across models, grouped by unguided CoT and guided CoT.

behavior is inconsistent. Finally, the results highlight the effectiveness of guided prompting. Most models successfully adhere to the provided instructions and generate the required reasoning steps. The primary exception is the Llama-3-1B model, which struggles to consistently execute the guided steps, suggesting a potential capability threshold for following complex procedural instructions.

5.2.2 QUALITY METRICS

The results of our manual evaluation for the quality metrics are presented in Table 1, where the quality of a given step is evaluated for the instances where the model actually performed that step.

The manual evaluation of model performance reveals several distinct patterns. Llama-3-1B is a clear outlier, struggling to produce complete and correct reasoning traces even when guided, which correlates with its low overall prediction accuracy. While both Llama-3-8B and -70B outperform their smaller counterpart, the relationship with scale is not linear; although the 70B model achieves a 15% higher overall score, it performs worse than the 8B model on several individual reasoning components. This indicates that larger model size does not necessarily guarantee a higher-quality reasoning trace. In contrast, Gemini-2.0-Flash benefits the most from a guided reasoning prompt, showing a 13% improvement in its overall score and significant gains across most components when guided. Generally, LLMs produce high-quality reasoning traces and final classifications when guided, and reasoning models achieve strong performance even in the unguided setting.

Finally, while the specialized reasoning models generally perform well, guided prompting offers them little benefit and, in some cases, even slightly degrades their performance. Notably, when provided with guided reasoning, the performance of non-reasoning models becomes comparable to that of the reasoning models. These findings suggest that guided non-reasoning models can offer equivalent quality to reasoning models. Given that our cost analysis (Table 9) shows that non-reasoning models are cheaper and more token-efficient, we suggest that using them in a guided scenario is more efficient than using unguided reasoning models.

Additionally, four general key findings emerge from this analysis:

- Guided prompting, while increasing the granularity of decompositions, reveals a critical trade-off between soundness and completeness. Asking for a more comprehensive breakdown improves completeness but often at the cost of soundness, as models generate subclaims not strictly entailed by the original statement. This suggests that the decomposition capabilities of current LLMs can be further improved.
- 2. Attribution correctness is high in the unguided reasoning setting, except for Gemini-2 and Llama-1B, though the scores are slightly lower in the guided setting. Additionally, both the inference and entailment steps achieve medium-to-high correctness (above 80% for most models) in both conditions.
- 3. The Aggregation step, when performed, is executed with almost-perfect accuracy across all models, meaning they almost consistently apply the correct aggregation logic.
- 4. With the notable exception of Llama-8B, we observe a consistent trend: for models where guided prompting enhanced the component-wise reasoning quality (e.g., Llama-1B, Gemini-2.0), there was a corresponding improvement in their overall prediction accuracy.

5.3 ATTRIBUTION TYPES

A deeper analysis of the attribution types reveals a clear hierarchy in their correctness. Extractive attribution, where the model directly quotes evidence from the source, is the most reliable, with an accuracy of 93%. Paraphrase attribution, where the model rephrases the source text, also demonstrates high correctness at 90%. In stark contrast, the accuracy of abstract attribution—which requires the model to provide a high-level summary of the relevant information—drops significantly to 69%.

5.4 Inter-Metric Correlations

To understand how the quality of one reasoning step relates to others, we computed the Spearman correlation between all quality metrics. We chose this coefficient because it is robust to outliers and non-linear relationships, making it well-suited for our medium-sized dataset. While we do not expect perfect correlations – as errors in one step can be corrected in a later one (e.g., in the Fig. 1 example, an error in one sub-claim's entailment decision does not prevent a correct final answer) – we hypothesized that a significant positive relationship exists.

The correlation matrix in Fig. 4 confirms this and reveals several key insights. The strongest correlations are between the later-stage components, with *Inference* and *Entailment* correctness being highly correlated with each other ($\rho=0.83$) and with the *Overall* score ($\rho=0.59$ and $\rho=0.62$, respectively). This suggests that proficiency in these steps may be the most critical driver of a correct answer. We also observe a strong correlation between Soundness and Completeness ($\rho=0.52$), indicating that models capable of producing sound decompositions also tend to produce complete ones. Conversely, Granularity shows a negligible correlation with all other metrics, reinforcing that simply creating more sub-claims is not an indicator of higher-quality reasoning traces. More correlation results are in Appendix C

5.5 LLM-AS-A-JUDGE

To scale our analysis, we employed <code>Gemini-2.5-Flash</code> as an LLM-as-a-Judge, which we first evaluated for agreement against our human-annotated judgments, calculating Gwet's score between the LLM's judgments and the manual judgments (full details are in Appendix D). The LLM usually demonstrated reliable performance, achieving 'Almost Perfect' or 'Substantial' agreement on 11 of our 13 key metrics, while only the inference and entailment metrics were rated 'Moderate'. This suggests that the LLM-as-a-Judge struggles to evaluate components where LLMs often produce errors in their reasoning process (as observed in Table 1). We then used this judge to evaluate the reasoning traces generated for the entire ClearFacts dataset, for the same models and prompts as in the manual analysis.

The full results of this large-scale analysis are presented in Appendix D. The results for the existence metrics (see Appendix D.2, Table 6) are consistent with the trends identified in our manual analysis (Section 5.2.1), confirming that models frequently omit reasoning steps unless explicitly guided, that larger models tend to spontaneously provide the NLI reasoning components, and that guided reasoning helps the model to actually perform those steps. Additionally, some of the results for the quality metrics (from Table 1) are observed in the automatic analysis of the entire dataset (full results in Appendix D.2, Table 7), while the ranking of the models performing the reasoning trace is similar to the ranking in the manual evaluation. Additionally, although the model is worse at evaluating the inference and entailment metrics (per the lower agreement with the manual evaluation for these two metrics), it still identifies that LLMs struggle to correctly provide those components, and it identifies that the better performing reasoning components are attribution and aggregation. This suggests that while there are phenomena we can observe using the LLM-as-a-Judge, there is still room for future work to improve the automatic evaluation of NLI reasoning.

6 Conclusion

In this work, we proposed a novel methodology for evaluating the reasoning traces of LLMs in the context of NLI reasoning for fact verification. Our framework moves beyond generic, step-by-step metrics by decomposing the NLI reasoning process into four distinct, functionally-motivated components: decomposition, attribution, inference, and aggregation, while evaluating the existence and quality of each component separately. Our manual and automated analyses of reasoning traces shed light on specific weaknesses in current models' reasoning, such as the frequent omission of reasoning steps and the trade-offs in quality that emerge with guided prompting. This diagnostic allows for a more targeted approach to improving model reliability.

Our findings open up several avenues for future work. The high-quality reasoning traces we identified can provide the basis for extracting concise justifications that will help users verify a model's final prediction, with a guided Gemini-2.0-Flash being a particularly cost-effective choice for this task. Additionally, these traces could be used as training data to distill the reasoning abilities of large models into smaller, more efficient ones like Llama-8B or Llama-70B. Finally, our evaluation approach itself can inspire the development of similar reasoning evaluation for other prominent NLP tasks, to enable a more fine-grained analysis of their reasoning processes.

REFERENCES

Fazl Barez, Tung-Yu Wu, Iván Arcuschin, Michael Lan, Vincent Wang, Noah Siegel, Nicolas Collignon, Clement Neo, Isabelle Lee, Alasdair Paren, Adel Bibi, Robert Trager, Damiano For-

- nasiere, John Yan, Yanai Elazar, and Yoshua Bengio. Chain-of-thought is not explainability, 2025.
 - Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Scott Wen tau Yih, and Yejin Choi. Abductive commonsense reasoning, 2020. URL https://arxiv.org/abs/1908.05739.
 - Siddhant Bhambri, Upasana Biswas, and Subbarao Kambhampati. Interpretable traces, unexpected outcomes: Investigating the disconnect in trace-based knowledge distillation, 2025. URL https://arxiv.org/abs/2505.13792.
 - Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. A large annotated corpus for learning natural language inference. In Lluís Màrquez, Chris Callison-Burch, and Jian Su (eds.), *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 632–642, Lisbon, Portugal, September 2015. Association for Computational Linguistics. doi: 10.18653/v1/D15-1075. URL https://aclanthology.org/D15-1075/.
 - Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. esnli: Natural language inference with natural language explanations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), Advances in Neural Information Processing Systems, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/4c7a167bb329bd92580a99ce422d6fa6-Paper.pdf.
 - Jiangjie Chen, Qiaoben Bao, Changzhi Sun, Xinbo Zhang, Jiaze Chen, Hao Zhou, Yanghua Xiao, and Lei Li. Loren: Logic-regularized reasoning for interpretable fact verification. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):10482–10491, June 2022. ISSN 2159-5399. doi: 10.1609/aaai.v36i10.21291. URL http://dx.doi.org/10.1609/aaai.v36i10.21291.
 - Jiaqi Chen, Bang Zhang, Ruotian Ma, Peisong Wang, Xiaodan Liang, Zhaopeng Tu, Xiaolong Li, and Kwan-Yee K. Wong. Spc: Evolving self-play critic via adversarial games for llm reasoning, 2025. URL https://arxiv.org/abs/2504.19162.
 - Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyou Zhou, and William Yang Wang. Tabfact: A large-scale dataset for table-based fact verification, 2020. URL https://arxiv.org/abs/1909.02164.
 - Ido Dagan, Oren Glickman, and Bernardo Magnini. The pascal recognising textual entailment challenge. In Joaquin Quiñonero-Candela, Ido Dagan, Bernardo Magnini, and Florence d'Alché Buc (eds.), *Machine Learning Challenges. Evaluating Predictive Uncertainty, Visual Object Classification, and Recognising Tectual Entailment*, pp. 177–190, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
 - Alvan R. Feinstein and Domenic V. Cicchetti. High agreement but low kappa: I. the problems of two paradoxes. *Journal of Clinical Epidemiology*, 43(6):543–549, 1990. ISSN 0895-4356. doi: https://doi.org/10.1016/0895-4356(90)90158-L. URL https://www.sciencedirect.com/science/article/pii/089543569090158L.
 - Kanishk Gandhi, Ayush Chakravarthy, Anikait Singh, Nathan Lile, and Noah D Goodman. Cognitive behaviors that enable self-improving reasoners, or, four habits of highly effective stars. *arXiv* preprint arXiv:2503.01307, 2025.
 - Olga Golovneva, Moya Peng Chen, Spencer Poff, Martin Corredor, Luke Zettlemoyer, Maryam Fazel-Zarandi, and Asli Celikyilmaz. ROSCOE: A suite of metrics for scoring step-by-step reasoning. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=xYlJRpzZtsY.
 - Kilem L Gwet. Handbook of inter-rater reliability: The definitive guide to measuring the extent of agreement among raters. Advanced Analytics, LLC, 2014.
 - Kilem Li Gwet. Computing inter-rater reliability and its variance in the presence of high agreement. *British Journal of Mathematical and Statistical Psychology*, 61(1):29–48, 2008.

Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenting Qi, Martin Riddell, Wenfei Zhou, James Coady, David Peng, Yujie Qiao, Luke Benson, Lucy Sun, Alexander Wardle-Solano, Hannah Szabó, Ekaterina Zubova, Matthew Burtell, Jonathan Fan, Yixin Liu, Brian Wong, Malcolm Sailor, Ansong Ni, Linyong Nan, Jungo Kasai, Tao Yu, Rui Zhang, Alexander Fabbri, Wojciech Maciej Kryscinski, Semih Yavuz, Ye Liu, Xi Victoria Lin, Shafiq Joty, Yingbo Zhou, Caiming Xiong, Rex Ying, Arman Cohan, and Dragomir Radev. FOLIO: Natural language reasoning with first-order logic. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 22017–22031, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.1229. URL https://aclanthology.org/2024.emnlp-main.1229/.

- Shibo Hao, Yi Gu, Haotian Luo, Tianyang Liu, Xiyan Shao, Xinyuan Wang, Shuhua Xie, Haodi Ma, Adithya Samavedhi, Qiyue Gao, Zhen Wang, and Zhiting Hu. LLM reasoners: New evaluation, library, and analysis of step-by-step reasoning with large language models. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=b0y6fbSUG0.
- Shreya Havaldar, Hamidreza Alvari, John Palowitch, Mohammad Javad Hosseini, Senaka Buthpitiya, and Alex Fabrikant. Entailed between the lines: Incorporating implication into NLI. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 32274–32290, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.1552. URL https://aclanthology.org/2025.acl-long.1552/.
- Hangfeng He, Hongming Zhang, and Dan Roth. SocREval: Large language models with the socratic method for reference-free reasoning evaluation. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Findings of the Association for Computational Linguistics: NAACL 2024*, pp. 2736–2764, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-naacl.175. URL https://aclanthology.org/2024.findings-naacl.175/.
- Subbarao Kambhampati, Kaya Stechly, Karthik Valmeekam, Lucas Saldyt, Siddhant Bhambri, Vardhan Palod, Atharva Gundawar, Soumya Rani Samineni, Durgesh Kalwar, and Upasana Biswas. Stop anthropomorphizing intermediate tokens as reasoning/thinking traces!, 2025. URL https://arxiv.org/abs/2504.09762.
- Ryo Kamoi, Tanya Goyal, Juan Diego Rodriguez, and Greg Durrett. Wice: Real-world entailment for claims in wikipedia. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 7561–7583, 2023.
- Xin Lai, Zhuotao Tian, Yukang Chen, Senqiao Yang, Xiangru Peng, and Jiaya Jia. Step-dpo: Step-wise preference optimization for long-chain reasoning of llms, 2024. URL https://arxiv.org/abs/2406.18629.
- J. Richard Landis and Gary G. Koch. The measurement of observer agreement for categorical data. Biometrics, 33(1):159–174, 1977. ISSN 0006341X, 15410420. URL http://www.jstor.org/stable/2529310.
- Deren Lei, Yaxi Li, Mengya Hu, Mingyu Wang, Vincent Yun, Emily Ching, and Eslam Kamal. Chain of natural language inference for reducing large language model ungrounded hallucinations, 2023. URL https://arxiv.org/abs/2310.03951.
- Miaoran Li, Baolin Peng, Michel Galley, Jianfeng Gao, and Zhu Zhang. Self-checker: Plug-and-play modules for fact-checking with large language models. In Kevin Duh, Helena Gomez, and Steven Bethard (eds.), *Findings of the Association for Computational Linguistics: NAACL 2024*, pp. 163–181, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-naacl.12. URL https://aclanthology.org/2024.findings-naacl.12/.
- Zhiyuan Li, Yi Chang, and Yuan Wu. Think-bench: Evaluating thinking efficiency and chain-of-thought quality of large reasoning models, 2025. URL https://arxiv.org/abs/2505.22113.

- Baohao Liao, Xinyi Chen, Sara Rajaee, Yuhui Xu, Christian Herold, Anders Søgaard, Maarten de Rijke, and Christof Monz. Lost at the beginning of reasoning, 2025. URL https://arxiv.org/abs/2506.22058.
 - Zimu Lu, Aojun Zhou, Ke Wang, Houxing Ren, Weikang Shi, Junting Pan, Mingjie Zhan, and Hongsheng Li. Step-controlled dpo: Leveraging stepwise error for enhanced mathematical reasoning, 2024. URL https://arxiv.org/abs/2407.00782.
 - Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. FActScore: Fine-grained atomic evaluation of factual precision in long form text generation. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, pp. 12076–12100, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.741. URL https://aclanthology.org/2023.emnlp-main.741/.
 - Abhika Mishra, Akari Asai, Vidhisha Balachandran, Yizhong Wang, Graham Neubig, Yulia Tsvetkov, and Hannaneh Hajishirzi. Fine-grained hallucination detection and editing for language models. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=dJMTn3QOWO.
 - Kushan Mitra, Dan Zhang, Sajjadur Rahman, and Estevam Hruschka. FactLens: Benchmarking fine-grained fact verification. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Findings of the Association for Computational Linguistics:* ACL 2025, pp. 18085–18096, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-256-5. doi: 10.18653/v1/2025.findings-acl.929. URL https://aclanthology.org/2025.findings-acl.929/.
 - Cheng Niu, Yuanhao Wu, Juno Zhu, Siliang Xu, Kashun Shum, Randy Zhong, Juntong Song, and Tong Zhang. Ragtruth: A hallucination corpus for developing trustworthy retrieval-augmented language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 10862–10878, 2024.
 - Md Rizwan Parvez. Chain of evidences and evidence to generate: Prompting for context grounded and retrieval augmented reasoning. In Weijia Shi, Wenhao Yu, Akari Asai, Meng Jiang, Greg Durrett, Hannaneh Hajishirzi, and Luke Zettlemoyer (eds.), *Proceedings of the 4th International Workshop on Knowledge-Augmented Methods for Natural Language Processing*, pp. 230–245, Albuquerque, New Mexico, USA, May 2025. Association for Computational Linguistics. ISBN 979-8-89176-229-9. doi: 10.18653/v1/2025.knowledgenlp-1.21. URL https://aclanthology.org/2025.knowledgenlp-1.21/.
 - Bibek Paudel, Alexander Lyzhov, Preetam Joshi, and Puneet Anand. Hallucinot: Hallucination detection through context and common knowledge verification, 2025. URL https://arxiv.org/abs/2504.07069.
 - Archiki Prasad, Swarnadeep Saha, Xiang Zhou, and Mohit Bansal. ReCEval: Evaluating reasoning chains via correctness and informativeness. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 10066–10086, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.622. URL https://aclanthology.org/2023.emnlp-main.622/.
 - Pengcheng Qiu, Chaoyi Wu, Shuyu Liu, Weike Zhao, Zhuoxia Chen, Hongfei Gu, Chuanjin Peng, Ya Zhang, Yanfeng Wang, and Weidi Xie. Quantifying the reasoning abilities of llms on real-world clinical cases, 2025. URL https://arxiv.org/abs/2503.04691.
 - Abulhair Saparov and He He. Language models are greedy reasoners: A systematic formal analysis of chain-of-thought. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=qFVVBzXxR2V.
 - Tal Schuster, Adam Fisch, and Regina Barzilay. Get your vitamin C! robust fact verification with contrastive evidence. In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek

Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 624–643, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.52. URL https://aclanthology.org/2021.naacl-main.52/.

- Wooseok Seo, Seungju Han, Jaehun Jung, Benjamin Newman, Seungwon Lim, Seungbeen Lee, Ximing Lu, Yejin Choi, and Youngjae Yu. Verifying the verifiers: Unveiling pitfalls and potentials in fact verifiers. *arXiv preprint arXiv:2506.13342*, 2025.
- Zayne Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. To cot or not to cot? chain-of-thought helps mainly on math and symbolic reasoning, 2025. URL https://arxiv.org/abs/2409.12183.
- Kaya Stechly, Karthik Valmeekam, Atharva Gundawar, Vardhan Palod, and Subbarao Kambhampati. Beyond semantics: The unreasonable effectiveness of reasonless intermediate tokens, 2025. URL https://arxiv.org/abs/2505.13775.
- Zhiqing Sun, Longhui Yu, Yikang Shen, Weiyang Liu, Yiming Yang, Sean Welleck, and Chuang Gan. Easy-to-hard generalization: Scalable alignment beyond human supervision. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=qwqfh2fTtN.
- Liyan Tang, Philippe Laban, and Greg Durrett. MiniCheck: Efficient fact-checking of LLMs on grounding documents. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 8818–8847, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.499. URL https://aclanthology.org/2024.emnlp-main.499/.
- Robert F. Tate. Correlation between a discrete and a continuous variable. point-biserial correlation. *The Annals of Mathematical Statistics*, 25(3):603–607, 1954. ISSN 00034851. URL http://www.jstor.org/stable/2236844.
- Onkar Thorat, Philippe Laban, and Chien-Sheng Wu. Summexecedit: A factual consistency benchmark in summarization with executable edits, 2025. URL https://arxiv.org/abs/2412.13378.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. FEVER: a large-scale dataset for fact extraction and VERification. In Marilyn Walker, Heng Ji, and Amanda Stent (eds.), *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 809–819, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1074. URL https://aclanthology.org/N18-1074/.
- Ran Tian, Shashi Narayan, Thibault Sellam, and Ankur P. Parikh. Sticking to the facts: Confident decoding for faithful data-to-text generation, 2020. URL https://arxiv.org/abs/1910.08684.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting, 2023. URL https://arxiv.org/abs/2305.04388.
- Andreas Vlachos and Sebastian Riedel. Fact checking: Task definition and dataset construction. In Cristian Danescu-Niculescu-Mizil, Jacob Eisenstein, Kathleen McKeown, and Noah A. Smith (eds.), *Proceedings of the ACL 2014 Workshop on Language Technologies and Computational Social Science*, pp. 18–22, Baltimore, MD, USA, June 2014. Association for Computational Linguistics. doi: 10.3115/v1/W14-2508. URL https://aclanthology.org/W14-2508/.
- Manya Wadhwa, Xinyu Zhao, Junyi Jessy Li, and Greg Durrett. Learning to refine with fine-grained natural language feedback. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen

(eds.), Findings of the Association for Computational Linguistics: EMNLP 2024, pp. 12281–12308, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.716. URL https://aclanthology.org/2024.findings-emnlp.716/.

- David Wan, Justin Chen, Elias Stengel-Eskin, and Mohit Bansal. MAMM-refine: A recipe for improving faithfulness in generation with multi-agent collaboration. In Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.), *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)*, pp. 9882–9901, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. ISBN 979-8-89176-189-6. doi: 10.18653/v1/2025.naacl-long.498. URL https://aclanthology.org/2025.naacl-long.498/.
- William Yang Wang. "liar, liar pants on fire": A new benchmark dataset for fake news detection. In Regina Barzilay and Min-Yen Kan (eds.), *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 422–426, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-2067. URL https://aclanthology.org/P17-2067/.
- Yingming Wang and Pepa Atanasova. Self-critique and refinement for faithful natural language explanations, 2025. URL https://arxiv.org/abs/2505.22823.
- Miriam Wanner, Seth Ebner, Zhengping Jiang, Mark Dredze, and Benjamin Van Durme. A closer look at claim decomposition. In Danushka Bollegala and Vered Shwartz (eds.), *Proceedings of the 13th Joint Conference on Lexical and Computational Semantics* (*SEM 2024), pp. 153–175, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.starsem-1.13. URL https://aclanthology.org/2024.starsem-1.13/.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems, 35:24824–24837, 2022.
- Xiaobao Wu, Liangming Pan, Yuxi Xie, Ruiwen Zhou, Shuai Zhao, Yubo Ma, Mingzhe Du, Rui Mao, Anh Tuan Luu, and William Yang Wang. AntiLeakBench: Preventing data contamination by automatically constructing benchmarks with updated real-world knowledge. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 18403–18419, Vienna, Austria, July 2025. Association for Computational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.901. URL https://aclanthology.org/2025.acl-long.901/.
- Fan Yang, Eduard Dragut, and Arjun Mukherjee. Claim verification under positive unlabeled learning. In *Proceedings of the 12th IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, ASONAM '20, pp. 143–150. IEEE Press, 2021. ISBN 9781728110561. doi: 10.1109/ASONAM49781.2020.9381336. URL https://doi.org/10.1109/ASONAM49781.2020.9381336.
- Barry Menglong Yao, Aditya Shah, Lichao Sun, Jin-Hee Cho, and Lifu Huang. End-to-end multimodal fact-checking and explanation generation: A challenging dataset and models. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '23, pp. 2733–2743, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450394086. doi: 10.1145/3539618.3591879. URL https://doi.org/10.1145/3539618.3591879.
- Fengzhu Zeng and Wei Gao. Prompt to be consistent is better than self-consistent? few-shot and zero-shot fact verification with pre-trained language models. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 4555–4569, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.278. URL https://aclanthology.org/2023.findings-acl.278/.

Yujun Zhou, Jiayi Ye, Zipeng Ling, Yufei Han, Yue Huang, Haomin Zhuang, Zhenwen Liang, Kehan Guo, Taicheng Guo, Xiangqi Wang, and Xiangliang Zhang. Dissecting logical reasoning in llms: A fine-grained evaluation and supervision study, 2025. URL https://arxiv.org/abs/2506.04810.

Ádám Kovács and Gábor Recski. Lettucedetect: A hallucination detection framework for rag applications, 2025. URL https://arxiv.org/abs/2502.17125.

A INTER-ANNOTATOR AGREEMENT (IAA)

To ensure the reliability of our annotations and establish a gold-standard evaluation set, we measured inter-annotator agreement (IAA). A subset of 30 randomly selected instances was independently annotated by all three annotators.

A key challenge of IAA measurement in this task is the frequent occurrence of highly imbalanced labels. For example, for the *Aggregation* step, nearly all instances include this component, leading to a skewed distribution. This phenomenon can cause traditional coefficients like Fleiss' Kappa to be misleadingly low despite high actual agreement (a.k.a the 'kappa paradox' (Feinstein & Cicchetti, 1990)).

To address this, we selected two complementary metrics:

- 1. Gwet's AC1/AC2 Coefficient: As our primary chance-corrected metric, we use Gwet's agreement coefficients, which are specifically designed to be robust in scenarios with skewed label distributions (Gwet, 2008; 2014). We use AC1 for binary metrics and AC2 for numerical metrics. This metric ranges from -1 to 1. We interpret the scores using the benchmarks from (Landis & Koch, 1977), where scores below 0.20 are considered 'poor', 0.21-0.40 'fair', 0.41-0.60 'moderate', 0.61-0.80 'substantial', and above 0.80 'almost perfect'.
- Raw Percent Agreement (RA): For maximum transparency, and similar to (Wu et al., 2025), we also report Raw Agreement. This intuitive metric measures the proportion of agreeing pairs among all possible rater pairs for a given item, averaged across all items.

Table 2: Inter-annotator agreement results, showing Gwet's score and Raw Percent Agreement (RA) values.

Metric	Existence		Corre	ctness				
	Gwet	RA	Gwet	RA				
Decomposition								
Decomposition	0.66	0.78	_	_				
Granularity (F1)	_	_	0.81	-				
Soundness	_	_	0.97	0.98				
Completeness	_	-	0.70	0.76				
Attrib	ution &	Entailm	ent					
Attribution	0.92	0.96	0.82	0.91				
Inference	0.96	0.96	0.87	0.90				
Entailment	0.84	0.87	0.75	0.81				
Aggregation								
Aggregation	0.92	0.93	0.87	0.89				
Overall Decision	_	_	0.89	0.91				

Since the metrics for Soundness, Attribution, and Entailment are evaluated on a per-sub-claim basis, and annotators might identify different sets of sub-claims from the same unstructured LLM output, we first established a consistent basis for comparison. We performed a pairwise alignment of the sub-claims extracted by each annotator and calculated the *F1-Score* to measure the consistency of their decompositions. This score, which balances precision and recall between the sets of extracted sub-claims (calculated as the number of sub-claims identified by both annotators, divided by the total number of sub-claims identified by annotator A; and vice versa), is reported as Granularity (F1) in Table 2. With this alignment established, we then calculated pairwise agreement for all other metrics and averaged the results across the three pairs of annotators (A-B, B-C, A-C).

The results demonstrate a high degree of reliability in our annotation scheme. Across the 12 core quality and existence metrics, the average chance-corrected agreement (Gwet's score) is 0.85 ('Almost Perfect'), and the average Raw Agreement is 0.89. A more detailed breakdown shows that 10

Table 3: Performance of the existence metrics of NLI reasoning across models annotated by the manual analysis. Each cell shows unguided CoT (left) vs. guided CoT (right). Highest values are in bold.

Model	Decomposition	Attribution	Inference	Entailment	Aggregation
Llama-3-1B	0.57 / 0.97	0.47 / 0.80	0.62 / 0.65	0.44 / 0.67	0.54 / 0.77
Llama-3-8B	0.93 / 1.00	0.83 / 0.97	0.89 / 0.92	0.67 / 0.94	0.88 / 0.93
Llama-3-70B	0.80 / 1.00	0.98 / 1.00	0.90 / 1.00	0.87 / 1.00	1.00 / 1.00
Gemini-2.0-Flash	0.30 / 0.90	0.97 / 1.00	0.83 / 0.98	0.87 / 0.98	1.00 / 1.00
DeepSeek-R1-distill Gemini-2.5-Flash	0.67 / 0.97 0.83 / 1.00	1.00 / 1.00 0.88 / 1.00	0.98 / 0.95 0.96 / 1.00	0.92 / 0.98 0.80 / 0.97	1.00 / 1.00 0.92 / 0.97

metrics achieved 'Almost Perfect' agreement, with the remaining 3 rated as 'Substantial', providing strong evidence for the validity of our data. The full results are available in Table 2.

We did observe that the lowest agreement scores were for the existence of the Decomposition and its Completeness. We hypothesize that this stems from the often unstructured and ambiguous nature of the LLM's output, which can make the identification of distinct sub-claims subjective. This was particularly true in the unguided CoT setting, where varied reasoning traces corresponded to lower IAA. In contrast, the uniform structures produced by guided prompts led to markedly higher agreement, though this trend was not statistically significant given our limited sample size.

B MANUAL ANALYSIS RESULTS

The full setup of the manual analysis is described in Section 5. The results for each model for each prompt, for the existence metrics are presented in Table 3. The results for the quality metrics are presented in Table 1.

C CORRELATION WITH FINAL PREDICTION

The correlation metrics between all the quality metrics are presented in Fig. 4 and described in Section 5.4.

To understand which reasoning steps correlate with the *overall* NLI decision, we conducted a correlation analysis. We measured the relationship between each metric's value and the final binary outcome (i.e., whether the model's overall prediction was correct). Given the binary nature of the outcome variable and the continuous nature of our metric scores, we used the *Point-Biserial correlation coefficient* $(r_{pb}; Tate (1954))$. In this setup, we investigate the relationship between the performance on each reasoning component and the final answer's correctness.

Table 4: Correlation of reasoning metrics with final prediction correctness. Values are point-biserial correlation coefficients (r_{pb}) with associated p-values. '(?)' indicates an existence metric, while '(#)' indicates a correctness metric. The two strongest correlates are highlighted in bold.

Metric	r_{pb}	p-value
Decomposition (?)	0.11	0.045
Granularity (#)	-0.06	0.283
Soundness (#)	0.21	8.0×10^{-5}
Completeness (#)	0.22	2.5×10^{-5}
Attribution (?)	0.18	4.6×10^{-4}
Attribution (#)	0.28	1.9×10^{-7}
Inference (?)	0.20	1.3×10^{-4}
Entailment (?)	0.25	1.3×10^{-6}
Inference (#)	0.44	2.7×10^{-16}
Entailment (#)	0.55	1.5×10^{-25}
Aggregation (?)	0.19	3.9×10^{-4}
Aggregation (#)	0.25	1.7×10^{-6}

The analysis, presented in Table 4, yields a crucial high-level insight: with the sole exception of *granularity*, all of our proposed reasoning metrics show a statistically significant, positive correlation with the final prediction's correctness. This demonstrates that the structured reasoning process, when performed correctly, is not merely superfluous text; it serves as a reliable witness to the validity of the final answer. The presence and quality of these intermediate steps are meaningfully linked to the model's success on the task.

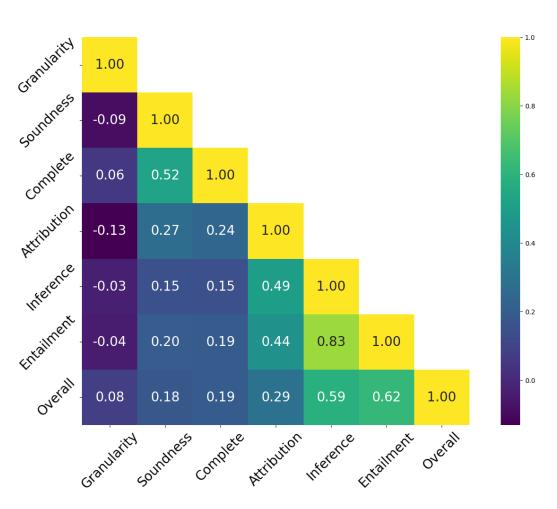


Figure 4: Spearman correlation matrix for the reasoning quality metrics. Yellow indicates a strong positive correlation, while dark purple indicates a weak or no correlation.

Table 6: LLM-as-a-Judge analysis results of the *existence metrics* across models. Each cell shows unguided CoT (left) vs. guided CoT (right), rounded to two decimals. Highest scores in **bold**.

Model	Decomposition	Attribution	Inference	Entailment	Aggregation
Llama-3-1B	0.22 / 0.90	0.62 / 0.62	0.57 / 0.63	0.57 / 0.73	0.80 / 0.66
Llama-3-8B	0.72 / 1.00	0.96 / 0.94	0.90 / 0.91	0.86 / 0.99	0.93 / 0.93
Llama-3-70B	0.85 / 1.00	0.99 / 0.98	0.94 / 0.94	0.94 / 1.00	1.00 / 1.00
Gemini-2.0-Flash	0.09 / 1.00	0.99 / 0.98	0.83 / 0.88	0.94 / 1.00	0.96 / 0.99
DeepSeek-R1-distill Gemini-2.5-Flash	0.24 / 0.93 0.44 / 1.00	0.99 / 0.98 0.97 / 0.99	0.87 / 0.93 0.87 / 0.90	0.97 / 0.99 0.92 / 1.00	0.99 / 1.00 0.98 / 1.00

Drilling deeper into these correlations, the results reveal a clear hierarchy of importance among the reasoning components. The strongest predictors of a correct final answer are, by a significant margin, the metrics evaluating the correctness of the entailment classification step. Specifically, entailment classification correctness shows the highest correlation ($r_{pb} = 0.55, p < 0.001$), followed by entailment inference correctness ($r_{pb} = 0.44, p < 0.001$). This finding underscores the intuitive conclusion that correctly solving the individual sub-problems is the most direct path to overall success on the NLI task.

Furthermore, the results highlight that the quality of a reasoning step is a more reliable indicator of success than its mere existence. For instance, attribution correctness ($r_{pb}=0.28$) has a considerably stronger correlation than attribution existence ($r_{pb}=0.18$). The only metric with no statistically significant correlation is granularity (p=0.28), indicating that the number of subclaims a model generates has no bearing on its success.

D LLM-AS-A-JUDGE

D.1 LLM-AS-A-JUDGE AGREEMENT

To further ground our insights of the quality of the reasoning steps of models in entailment decision, we use LLM-as-a-Judge to evaluate the reasoning trace of the whole dataset. We instruct a model to produce all the existence and quality metrics, similar to the annotators. Out of 13 metrics, 1 metric is 'Almost Perfect', 10 are 'Substantial', and 2 are in the high range of 'Moderate'. The full results are in Table 5.

Table 5: Agreement between LLM-as-a-Judge and the evaluation set. Each cell shows the Judge value.

Metric	Existence		Qua	lity			
	Gwet	RA	Gwet	RA			
Decomposition							
Decomposition	0.75	0.84	_	_			
Granularity	_	_	0.86	0.87			
Soundness	_	_	0.67	0.71			
Completeness	_	-	0.81	0.85			
Attribution & Entailment							
Attribution	0.78	0.80	0.63	0.67			
Inference	0.72	0.75	0.56	0.61			
Entailment	0.75	0.77	0.58	0.63			
Aggregation							
Aggregation	0.89	0.90	0.77	0.83			
Overall Decision	_	_	0.60	0.75			

D.2 RESULTS

The full LLM-as-a-Judge experiment setup is described in Section 5.5. The results on the *existence* metrics are in Table 6. The results on the *quality* metrics are presented in Section 5. The differences in the observations between the manual and the automated evaluations are described in Section 5.5.

When we group the agreement data by prompt type, a gap in a few metrics becomes distinguishable. The agreement on the existence of decomposition in the unguided CoT setting was low (a Gwet's score of 0.41), compared to a near-perfect score of 0.98 in the guided setting. A similar gap was observed for granularity, with a Gwet's score of 0.75 for unguided prompts versus 0.96 for guided prompts. The full results are in Table 8.

Table 7: LLM-as-a-Judge analysis results on the *correctness metrics*. Each cell shows unguided CoT (left) vs. guided CoT (right), rounded to two decimals. Highest values in each column are in **bold**.

Model	Granularity	Soundness	Completeness	Attribution
Llama-3-1B	2.89 / 3.46	0.70 / 0.45	0.60 / 0.30	0.63 / 0.54
Llama-3-8B	2.57 / 2.96	0.94 / 0.82	0.93 / 0.78	0.89 / 0.80
Llama-3-70B	3.01 / 2.95	0.97 / 0.93	0.96 / 0.95	0.97 / 0.94
Gemini-2.0-Flash	2.71 / 2.37	0.96 / 0.94	0.96 / 0.97	0.97 / 0.94
DeepSeek-R1-distill	3.04 / 2.96	0.98 / 0.96	0.98 / 0.99	0.97 / 0.95
Gemini-2.5-Flash	2.89 / 3.35	0.96 / 0.98	0.98 / 0.99	0.98 / 0.98
Model	Inference	Entailment	Aggregation	Overall
-				

Model	Inference	Entailment	Aggregation	Overall
Llama-3-1B	0.46 / 0.41	0.46 / 0.35	0.38 / 0.28	0.53 / 0.54
Llama-3-8B	0.77 / 0.71	0.77 / 0.69	0.70 / 0.66	0.65 / 0.68
Llama-3-70B	0.92 / 0.90	0.91 / 0.89	0.90 / 0.93	0.84 / 0.85
Gemini-2.0-Flash	0.89 / 0.91	0.87 / 0.90	0.92 / 0.90	0.82 / 0.85
DeepSeek-R1-distill	0.92 / 0.91	0.90 / 0.90	0.91 / 0.96	0.84 / 0.84
Gemini-2.5-Flash	0.93 / 0.97	0.91 / 0.95	0.94 / 0.98	0.87 / 0.86

Table 8: Comparison between unguided and guided CoT prompts for the agreement between the LLM-as-a-Judge and the human annotations. AC = Gwet's AC1, RA = Raw Agreement.

Metric		Existence Cor			Corre	rectness		
Nicorie	Ung	uided	Gui	ided	Ung	Unguided		ded
	AC	RA	AC	RA	AC	RA	AC	RA
Decomposition								
Decomposition	0.41	0.70	0.98	0.98	l –	_	-	_
Atomicity	_	_	_	_	0.75	0.78	0.96	0.97
Soundness	_	_	_	_	0.62	0.67	0.71	0.74
Completeness	-	-	_	_	0.84	0.87	0.78	0.83
		Attribu	tion &	Entailm	ent			
Attribution	0.67	0.71	0.87	0.88	0.60	0.65	0.65	0.69
Inference	0.60	0.65	0.83	0.84	0.54	0.60	0.57	0.62
Entailment	0.61	0.66	0.87	0.88	0.55	0.60	0.61	0.65
Aggregation								
Aggregation	0.83	0.86	0.94	0.95	0.76	0.82	0.78	0.84
Overall Decision	_	_	_	_	0.65	0.79	0.55	0.72

E NUMBER OF TOKENS AND COST

An analysis of the computational cost, presented in Table 9, reveals important efficiency trade-offs. Excluding Llama-1B, we find that non-reasoning models generally produce fewer output tokens than their reasoning-focused counterparts, even in the guided scenario. However, the average cost for a guided non-reasoning model is substantially lower than for an unguided reasoning model. Given that the quality of the reasoning and the overall performance are equivalent between these two configurations (as established in Section 5.2), our findings suggest that using guided prompting with non-reasoning models is a more efficient and cost-effective strategy.

Table 9: Average output tokens and cost across models. Results are reported for unguided CoT and guided CoT. Costs are based on the current price of the models in the OpenRouter supplier, in US dollars.

Model	Unguid	ed CoT	Guided CoT		
1120 401	Output Tokens Cost ($\times 10^{-3}$ \$)		Output Tokens	Cost ($\times 10^{-3}$ \$)	
Llama-3-1B	312.7	0.003	844.5	0.008	
Llama-3-8B	243.9	0.007	346.9	0.010	
Llamam-3-70B	285.5	0.010	329.3	0.012	
Gemini-2.0-Flash	121.6	0.049	289.0	0.116	
DeepSeek-R1-distill	443.9	0.120	548.8	0.148	
Gemini-2.5-Flash	247.4	0.619	1398.1	3.495	

F PROMPTS

We evaluated all models using two distinct prompt conditions. The *unguided CoT* condition utilized a standard chain-of-thought prompt, taken from (Wan et al., 2025), which encourages the model to produce a free-form reasoning trace. The *guided CoT* condition augmented this prompt with explicit instructions, directing the model to adhere to the specific, systematic reasoning structure defined in Section 3. The unguided CoT prompt is Prompt no. 1, and the guided CoT prompt is Prompt no. 2.

Prompt 1: NLI Unguided CoT

Document: {document}

Sentence: {claim}

Determine if the sentence is factually consistent with the document provided above. A sentence is factually consistent if it can be entailed (either stated or implied) by the document. If any part of the claim is not substantiated, it should be considered inconsistent.

Let's think step by step.

Conclude your response with either "yes" (the claim is consistent) or "no" (the claim is inconsistent).

Prompt 2: NLI Guided CoT Document: {document} Sentence: {claim} Steps: - For each sub-claim: relevant information. 2. Classify the sub-claim as: - Finally, decide: sistent).

Determine if the sentence is factually consistent with the document provided above. tence is factually consistent if it can be entailed (either stated or implied) by the document. If any part of the claim is not substantiated, it should be considered inconsistent.

- Decompose the claim into distinct sub-claims.
- 1. Identify the exact text in the document that supports or contradicts it, or note that there is no
 - - entailed fully supported or implied by the document
 - contradicted directly refuted by the document
 - neutral neither supported nor contradicted (no evidence)
- Provide your reasoning by listing each sub-claim with its classification and evidence.
 - · "yes" if all sub-claims are entailed
 - "no" if any sub-claim is contradicted or neutral

Conclude your response with either "yes" (the claim is consistent) or "no" (the claim is incon-

LLM USAGE

Throughout the writing process of this paper, we utilized LLMs to assist with polishing the text, including correcting grammatical errors and improving clarity through paraphrasing.