

# Towards Effective Extraction and Evaluation of Factual Claims

Dasha Metropolitansky, Jonathan Larson

Microsoft Research

{dasham, jolarso}@microsoft.com

## Abstract

A common strategy for fact-checking long-form content generated by Large Language Models (LLMs) is extracting simple claims that can be verified independently. Since inaccurate or incomplete claims compromise fact-checking results, ensuring claim quality is critical. However, the lack of a standardized evaluation framework impedes assessment and comparison of claim extraction methods. To address this gap, we propose a framework for evaluating claim extraction in the context of fact-checking along with automated, scalable, and replicable methods for applying this framework, including novel approaches for measuring coverage and decontextualization. We also introduce Claimify, an LLM-based claim extraction method, and demonstrate that it outperforms existing methods under our evaluation framework. A key feature of Claimify is its ability to handle ambiguity and extract claims only when there is high confidence in the correct interpretation of the source text.

## 1 Introduction

It is well known that Large Language Models (LLMs) are prone to producing unsubstantiated or inaccurate content (Huang et al., 2025). As LLM-generated content grows in volume and influence, reliable fact-checking systems become increasingly important.

For long-form, information rich outputs, a common fact-checking strategy is to extract simple “claims” from the text, then retrieve relevant evidence and assess the veracity of each claim independently (Min et al., 2023; Hu et al., 2024a). The effectiveness of such “decompose-then-verify” systems is contingent on the quality of the extracted claims: misrepresenting the source text or omitting factual content can result in misleading or incomplete conclusions. Therefore, rigorous evaluation of claim extraction methods is critical.

While prior works have identified desirable properties of claims, classified common errors, and shown that fact-checking performance is sensitive to the decomposition method, there is currently no standardized approach for evaluating claim extraction (Hu et al., 2024a; Wanner et al., 2024b).

This paper makes the following contributions:

1. We propose a framework for evaluating claim extraction methods in the context of fact-checking. We also introduce automated, scalable, and replicable methods for applying this framework, which are validated through human review. Two key innovations are: (1) a granular assessment of claims’ coverage of the source text, and (2) an outcome-based approach for evaluating decontextualization (i.e., whether a claim contains all necessary contextual information).
2. We introduce Claimify, an LLM-based claim extraction method. We demonstrate that it outperforms existing methods under our evaluation framework. To the best of our knowledge, Claimify is also the first claim extraction method that identifies sentences with multiple possible interpretations and determines when the correct interpretation cannot be inferred from the sentence’s context – unlike existing methods, which either ignore ambiguity or assume it is always resolvable.

## 2 Evaluating Claim Extraction

### 2.1 Key Concepts

The definition of a “claim” varies across prior works (Daxenberger et al., 2017). We adopt the perspective from Ni et al. (2024), which focuses on statements that “present verifiable facts,” where a fact is “a statement or assertion that can be objectively verified as true or false based on empirical evidence or reality.” We use the term “factual claim”

throughout this paper instead of the full phrase “verifiable factual claim.”

We argue that in the context of fact-checking, claim extraction methods should be evaluated based on three factors:

1. **Entailment** means that if the source text is true, the extracted claims must also be true. The broader principle that the source text should support the claims has been described in previous works as faithfulness (Wright et al., 2022; Hu et al., 2024b; Chen et al., 2024), coherence (Wanner et al., 2024b), and correctness (Kamoi et al., 2023).
2. **Coverage** means that extracted claims should capture the verifiable information in the source text without explicitly including the unverifiable information. We discuss a novel approach to evaluating coverage in §2.2.
3. **Decontextualization** is typically defined as: (1) each claim should be understandable on its own, without requiring additional context, and (2) each claim should retain the meaning it held in its original context (Choi et al., 2021; Gunjal and Durrett, 2024). We propose an alternative definition in §2.3.

Claim extraction methods have also been evaluated based on **atomicity** (Wanner et al., 2024b; Chen et al., 2024). For example, the claim “California and New York implemented a plastic bag ban” is not atomic because it can be divided into “California implemented a plastic bag ban” and “New York implemented a plastic bag ban.” However, the pursuit of atomicity lacks a clear endpoint: the above claims could be further divided into “At least one state has implemented a plastic bag ban,” “California has implemented a ban,” and “California exists.” Moreover, prior works suggest that atomicity does not consistently improve fact-checking performance (Chen et al., 2023a; Hu et al., 2024a; Tang et al., 2024). As a result, we do not consider atomicity in our evaluation framework.

## 2.2 Coverage

Coverage (as defined in §2.1) can be evaluated at different levels of granularity. Prior works have primarily focused on **sentence-level** evaluation, assessing whether a method correctly determines that a sentence, as a whole, contains a factual claim (Konstantinovskiy et al., 2021; Majer and Šnajder, 2024).

Consider the sentence “*The iconic American flag has 50 stars and 13 stripes,*” where Method A extracts the claims [“*The American flag is iconic*”, “*The American flag has numerous stars and stripes*”] and Method B extracts the claims [“*The American flag contains 50 stars*”, “*The American flag contains 13 stripes*”]. Both methods correctly identified that the sentence contains a factual claim, so they performed equally well in terms of sentence-level coverage.

In contrast, we introduce the concept of **element-level coverage**, evaluated by breaking a sentence into distinct pieces of information (“elements”), classifying each element as verifiable or unverifiable, then assessing whether each element is “covered” by (i.e., present in) the extracted claims, either explicitly or implicitly (i.e., the element is stated or suggested).

We define a **true positive** as a verifiable element that is covered implicitly or explicitly by the claims; a **true negative** as an unverifiable element that is either not covered or only implicitly covered (since implicit coverage may not reflect deliberate inclusion); a **false positive** as an unverifiable element that is explicitly covered; and a **false negative** as a verifiable element that is not covered. In Appendix A, we provide an example that illustrates why explicit coverage of unverifiable elements should count as a false positive, but implicit coverage should not.

Unlike sentence-level coverage, element-level coverage recognizes that Method B is superior to Method A. The sentence contains three elements: (1) the American flag is iconic, (2) the American flag has 50 stars, and (3) the American flag has 13 stripes. Only elements 2 and 3 are verifiable. Method A has one false positive (it explicitly covers element 1) and two false negatives (it does not mention the number of stars and stripes, so it fails to cover elements 2 and 3), while Method B has one true negative (it does not cover element 1) and two true positives (it covers elements 2 and 3).

Prior works that came closest to evaluating element-level coverage, such as Song et al. (2024) and Li et al. (2024), (1) relied on human annotation, making them difficult to scale, (2) lacked specificity (e.g., they considered whether verifiable content was omitted without quantifying the omissions), (3) failed to penalize the inclusion of unverifiable content, and/or (4) did not distinguish between implicit and explicit coverage.

### 2.3 Decontextualization

Numerous studies rely on human annotations to assess whether a unit of text (e.g., a sentence or claim) is sufficiently decontextualized (Choi et al., 2021; Kane and Schubert, 2023; Bayat et al., 2025). However, we argue that such judgments are often subjective, difficult to apply consistently, and fail to reflect the claim’s suitability for fact-checking.

Consider the claim “*John Smith supports government regulations*” extracted from the sentence “*In the latest episode of Jane Doe’s podcast on electric vehicles, Doe’s free-market views clashed with John Smith’s support for government regulations.*” According to the definition in §2.1, the claim appears sufficiently decontextualized.

However, if a fact-checking system attempted to verify this claim, it might find evidence of John Smith opposing government regulations in other contexts (e.g., AI or healthcare) and conclude that the claim is false – even though this evidence does not contradict the source sentence. The mismatch between the evidence’s implications for the claim and for the sentence indicates that the claim was insufficiently decontextualized: it should have clarified that John’s comments were made during a specific podcast episode about electric vehicles. Critically, the underspecification only became apparent after the fact-checking process, not beforehand.

Consider a second example: “*The court helped secure Bush’s presidency through its split decision to halt the Florida recount.*” The claim appears to be insufficiently decontextualized, since “*The court*” is not defined. However, for fact-checking purposes, the underspecification is not problematic, since the only plausible reference is the Bush v. Gore decision by the United States Supreme Court.

We posit that instead of making subjective judgments about whether a claim is “sufficiently” decontextualized, we should measure how the claim affects the outcome of the fact-checking system. In a fact-checking system, claims are used to retrieve evidence from a collection of documents, which informs a true/false verdict. Therefore, missing context is problematic only if its inclusion would change the verdict from true to false, or vice versa.<sup>1</sup> This shift could occur if including the context results in retrieving a different pool of evidence with the opposite relationship to the claim, or if the same

<sup>1</sup>Prior works propose retrieval-based evaluations of decontextualization (Choi et al., 2021; Deng et al., 2024). However, in fact-checking, such approaches are insufficient because retrieval is only an intermediate step towards the final verdict.

evidence is retrieved but its relationship to the claim changes when viewed with the added context.

Accordingly, we propose a three-step process for evaluating the decontextualization of a claim  $c$  in the context of fact-checking:

1. **Identify Missing Context.** Based on  $c$  and its context, either:

- Generate  $c_{\max}$ , a maximally decontextualized version of  $c$ , ensuring  $c$  is entailed by  $c_{\max}$ ; or
- Determine that  $c$  is already maximally decontextualized (i.e.,  $c = c_{\max}$ )

In the John Smith example,  $c_{\max}$  might be: “*In the latest episode of Jane Doe’s podcast on electric vehicles, John Smith supports government regulations.*” If  $c$  is already maximally decontextualized, no further steps are needed.<sup>2</sup>

2. **Retrieve Evidence.** In the collection of documents used for fact-checking, find relevant information for  $c$  and  $c_{\max}$ , producing evidence sets  $E_c$  and  $E_{\max}$ , respectively.

3. **Determine Veracity.** Perform the following checks<sup>3</sup>:

- $E_c \Rightarrow c$  (i.e., check if  $E_c$  supports  $c$ )
- $E_{\max} \Rightarrow c_{\max}$
- If  $E_c \Rightarrow c$ , check if  $E_c \Rightarrow c_{\max}$

This process yields one of seven possible results:

1.  $c = c_{\max}$
2.  $(E_c \Rightarrow c) \wedge (E_{\max} \Rightarrow c_{\max}) \wedge (E_c \Rightarrow c_{\max})$
3.  $(E_c \Rightarrow c) \wedge (E_{\max} \Rightarrow c_{\max}) \wedge (E_c \not\Rightarrow c_{\max})$
4.  $(E_c \Rightarrow c) \wedge (E_{\max} \not\Rightarrow c_{\max}) \wedge (E_c \Rightarrow c_{\max})$
5.  $(E_c \Rightarrow c) \wedge (E_{\max} \not\Rightarrow c_{\max}) \wedge (E_c \not\Rightarrow c_{\max})$
6.  $(E_c \not\Rightarrow c) \wedge (E_{\max} \Rightarrow c_{\max})$
7.  $(E_c \not\Rightarrow c) \wedge (E_{\max} \not\Rightarrow c_{\max})$

<sup>2</sup>Just as there are often multiple ways to decontextualize a sentence, there is rarely a single “correct” formulation of  $c_{\max}$ . However, we argue that creating a claim that contains as much context as possible is less subjective than trying to evaluate whether a claim is “sufficiently” decontextualized. The evaluation can also be repeated for different  $c_{\max}$  values to ensure robustness.

<sup>3</sup>If  $E_c \not\Rightarrow c$ , there is no need to check whether  $E_c \Rightarrow c_{\max}$ . Since  $c_{\max}$  entails  $c$ , and  $c$  is narrower than  $c_{\max}$ , any evidence that fails to support  $c$  cannot support  $c_{\max}$ .

Results 5 and 6 are undesirable because the verdicts for  $c$  and  $c_{\max}$  are misaligned. Result 3 – where  $c$  and  $c_{\max}$  are supported by their respective evidence sets, but the evidence for  $c$  does not support  $c_{\max}$  – is problematic in scenarios where the rationale matters, not just the verdict.<sup>4</sup>

In contrast, Results 2 and 7 are desirable because the verdicts for  $c$  and  $c_{\max}$  are aligned. Result 4 is also favorable – in fact, it suggests that  $c$  is superior to  $c_{\max}$  because only the former retrieved evidence supporting both  $c$  and  $c_{\max}$ . We classify Result 1 as desirable because it indicates that no contextual information was omitted.

We describe our implementation of this approach in §5.3, comparing claim extraction methods based on the percentage of desirable results.

### 3 Claimify

This section describes Claimify, our novel LLM-based claim extraction method. Figure 1 in Appendix B illustrates its key stages, and Appendix N.1 contains all prompts.

#### 3.1 Sentence Splitting and Context Creation

Claimify accepts a question-answer pair as input. It uses NLTK’s sentence tokenizer to split the answer into sentences (Bird and Loper, 2004, version 3.9.1). Context is created for each sentence  $s$  based on a configurable combination of  $p$  preceding sentences,  $f$  following sentences, and optional metadata (e.g., the header hierarchy in a Markdown-style answer).<sup>5</sup> The parameters  $p$  and  $f$  are defined separately for the stages outlined in §3.2–§3.4, allowing each stage to have a distinct context.

#### 3.2 Selection

Next, Claimify uses an LLM to determine whether each sentence contains any verifiable content, in light of its context and the question. When the LLM identifies that a sentence contains both verifiable and unverifiable components, it rewrites the

sentence, retaining only the verifiable components.

More specifically, the LLM selects one of the following options: (1) state that the sentence does not contain any verifiable content, (2) return a modified version of the sentence that retains only verifiable content, or (3) return the original sentence, indicating that it does not contain any unverifiable content. If the LLM selects the first option, the sentence is labeled “No verifiable claims” and excluded from subsequent stages (§3.3 and §3.4). Table 5 in Appendix B provides examples where the LLM selected the first or second option.

#### 3.3 Disambiguation

The primary goals of this stage are to identify ambiguity in the sentences returned by the Selection stage, and to determine whether the ambiguity has a clear resolution based on the question and the context. These objectives and capabilities are unique to Claimify (see §7 for a discussion of related works).

Claimify uses an LLM to identify two types of ambiguity. The first is **referential ambiguity**, which occurs when it is unclear what a word or phrase refers to. For example, in the sentence “*They will update the policy next year;*” the terms “*They*,” “*the policy*,” and “*next year*” are ambiguous. The second is **structural ambiguity**, which occurs when grammatical structure allows for multiple interpretations. For instance, the sentence “*AI has advanced renewable energy and sustainable agriculture at Company A and Company B*” can be interpreted as: (1) AI has advanced renewable energy and sustainable agriculture at both Company A and Company B, or (2) AI has advanced renewable energy at Company A, and it has advanced sustainable agriculture at Company B.

A special case of structural ambiguity involves distinguishing between factual claims and unverifiable interpretations added by the author. For example, the sentence “*John emphasized the support he received from executives throughout his career; highlighting the importance of mentorship,*” can be interpreted as: (1) John both emphasized the support he received and highlighted the importance of mentorship, or (2) John emphasized the support he received, while the author added the interpretation about the importance of mentorship.

The LLM is also asked to determine whether each instance of ambiguity can be resolved using the question and the context. The standard for resolution is whether a group of readers would likely agree on the correct interpretation. For example,

<sup>4</sup>Consider  $c =$  “*Miller has been described as an architect,*” extracted from the sentence  $s =$  “*Miller has been described as the architect of Trump’s controversial immigration policies.*” Let  $c_{\max} = s$ . Imagine  $E_c$  contains information about a building architect named John Miller, while  $E_{\max}$  describes Stephen Miller, President Trump’s policy advisor, as an architect of the administration’s immigration policies. Although both  $c$  and  $c_{\max}$  are supported by their respective evidence sets, it would be highly problematic if a fact-checking system cited  $E_c$  as its rationale for the sentence’s veracity! Note that  $c$  and  $s$  were adapted from an example by Wanner et al. (2024a).

<sup>5</sup>We did not use any metadata for the experiments described in this paper.

recall the sentence “*AI has advanced renewable energy and sustainable agriculture at Company A and Company B.*” If the context specified that Company A builds solar panels and Company B reduces farms’ water usage, readers would likely conclude that AI has advanced renewable energy at Company A and sustainable agriculture at Company B. Conversely, if the context only described both companies as “*environmental pioneers,*” readers would have insufficient information to determine the correct interpretation.

If any ambiguity is unresolvable, the sentence is labeled “Cannot be disambiguated” and excluded from the Decomposition stage (§3.4), even if it has unambiguous, verifiable components. Table 6 in Appendix B provides examples of such sentences. If the LLM resolves all ambiguity, it returns a clarified version of the sentence. If there is no ambiguity, it returns the original sentence.<sup>6</sup>

Across the models tested in our experiments (§5), the largest proportion of sentences labeled “Cannot be disambiguated” was 5.4%.<sup>7</sup>

### 3.4 Decomposition

In the final stage, Claimify uses an LLM to decompose each disambiguated sentence into decontextualized factual claims. If it does not return any claims (only 0.8% of cases in our experiments), the sentence is labeled “No verifiable claims.”

Extracted claims may include text in brackets, which typically represents information implied by the question or context but not explicitly stated in the source sentence. For example, given the question “*Provide an overview of celebrities’ stances on the Middle East,*” and the sentence “*John has called for peace,*” Claimify may return the claim “*John [a celebrity] has called for peace [in the Middle East].*” This notation resembles the “markup-and-mask” approach by Eisenstein et al. (2024), which adds bracketed text to clarify context in passages. A benefit of bracketing is that it flags inferred content, which is inherently less reliable than content explicitly stated in the source sentence.

<sup>6</sup>The LLM also checks for partial names, abbreviations, and acronyms, which are not considered linguistic ambiguities. If full forms are provided in the question or context, the LLM includes them in the returned sentence; otherwise, the LLM leaves them unchanged to avoid factual inaccuracies.

<sup>7</sup>The proportion of sentences labeled “Cannot be disambiguated” per model was as follows: mistral-large-2411 = 5.4%, gpt-4o-2024-08-06 = 3.2%, DeepSeek-V3 = 2.4%.

## 4 Experimental Setup

### 4.1 Data

#### 4.1.1 BingCheck

We evaluated Claimify’s performance on the BingCheck dataset (Li et al., 2024), which consists of 396 answers generated by Microsoft Copilot (formerly Bing Chat). BingCheck spans a wide range of topics and question types, and its answers are significantly longer than those in comparable datasets (Li et al., 2024). As a result, it reflects the diversity and complexity of real-world LLM usage in long-form question answering. Moreover, since BingCheck answers are generated based on web search results, it is reasonable to expect that relevant evidence exists for many claims – a key consideration for the evidence retrieval step of the decontextualization evaluation described in §2.3.

#### 4.1.2 Human Annotation Study

We conducted a human annotation study to classify sentences in BingCheck answers as containing or not containing factual claims.<sup>8</sup> A total of 6,490 sentences were labeled by three annotators who are familiar with natural language processing, including one of the authors. To ensure reliability, annotators completed two practice rounds on a subset of sentences, resolving disagreements via discussion, then independently annotated the remaining data. Krippendorff’s alpha (Krippendorff, 2013; Castro, 2017) increased from 0.44 in the first practice round to 0.72 in the final round, reaching 0.86 for high-confidence annotations. Sentence splitting methodology, annotation procedure, guidelines, and results are detailed in Appendix C. The labels from the study were used in our analysis of coverage (§5.2).

### 4.2 Baseline Methods

We compared Claimify to five LLM-based methods:

1. **AFaCTA (Ni et al., 2024)** uses an ensemble of prompts to classify sentences as containing or not containing objectively verifiable content.
2. **Factcheck-GPT (Wang et al., 2024)** classifies sentences as factual claims, opinions, non-claims (e.g., questions or imperative statements), or other.

<sup>8</sup>The dataset will be released at <https://aka.ms/claimify-dataset>.

3. **VeriScore** (Song et al., 2024) combines sentence classification, decomposition, and decontextualization in a single prompt. It returns either “No verifiable claim” or a list of claims.
4. **DnD** (Wanner et al., 2024a) decomposes and decontextualizes sentences in a single prompt.
5. **SAFE** (Wei et al., 2024) adds instructions to FActScore’s decomposition prompt (Min et al., 2023) and performs decontextualization in a separate prompt.

DnD and SAFE do not provide instructions for handling sentences without factual claims. Therefore, when the LLM declined to extract claims, it did not use a consistent output format. If no claims were parsed from the output, we assumed that the LLM determined there were no factual claims.

We selected these methods because they allow for direct comparisons with Claimify. They process sentences independently, unlike other methods that analyze the entire answer as a single unit (e.g., Chern et al., 2023; Bayat et al., 2025). The methods with explicit sentence classification components (AFaCTA, Factcheck-GPT, VeriScore) share Claimify’s focus on detecting verifiable content, rather than ranking sentences by their “check-worthiness” (see §7). The methods that perform claim extraction (VeriScore, DnD, SAFE) involve both decomposition and decontextualization, unlike other approaches that focus solely on decomposition (e.g., Kamoi et al., 2023; Chen et al., 2023a).

To further enable direct comparisons, we used the sentence splitting logic described in Appendix C.1 for all methods. We also made minimal edits to all prompts (except VeriScore, where edits were unnecessary) to include the question and clarify that the sentence was extracted from a response to the question. Additional settings are described in Appendix D.

The claim extraction methods (VeriScore, DnD, SAFE, Claimify) generated a total of 73,681 claims. Where a method produced duplicate claims for a sentence, we removed the duplicates, resulting in 73,229 claims. All subsequent sections refer to this de-duplicated claim set.

## 5 Experiments

This section describes our implementation of the evaluation framework outlined in §2 and the corresponding results. Appendix N.2 contains all prompts. Appendix E provides the sentence context

definitions, and Appendix F describes the samples for all experiments. All results in this section were produced using OpenAI’s gpt-4o-2024-08-06 model with a temperature of 0; results for other models are reported in Appendix G. All reported p-values were derived from two-proportion Z-tests, with Holm-Bonferroni correction for multiple comparisons.

### 5.1 Entailment

To determine whether claims are entailed by their source sentences, we first used a pre-trained Natural Language Inference (NLI) model from Nie et al. (2020), as was done by Wanner et al. (2024b). We tried two configurations, both of which revealed significant limitations, detailed in Appendix H.

In light of the NLI model’s limitations, we developed a prompt that classifies a claim as entailed or not entailed based on the source sentence, context, and question. To validate the prompt, we randomly sampled 20 claims from each claim extraction method (80 claims total) and labeled them without referencing the LLM’s outputs. The LLM’s classifications conflicted with our labels in only five cases, whereas the NLI model’s classifications conflicted with us in 32 and 12 cases for the first and second configurations, respectively. Appendix I provides an overview of cases where we disagreed.

Table 1 shows the percentage of entailed claims for each method using our prompt. Claimify and VeriScore achieved the highest percentage of entailed claims (99%), with no statistically significant difference between them ( $p=0.145$ ). All pairwise comparisons between the methods, except for Claimify vs. VeriScore, showed statistically significant differences ( $p<0.001$ ). These results align with a similar analysis by Wanner et al. (2024b) where the percentage of supported claims for various claim extraction methods, averaged across different models, ranged from 86% to 98%.

Method	Claims	% Entailed
Claimify	12,406	99.0
DnD	27,717	89.1
SAFE	22,786	96.6
VeriScore	7,420	99.2

Table 1: Percentage of claims entailed by the combined source sentence, context, and question, along with the total number of claims (as described in Appendix F.2), per method.

Method	Accuracy		Macro $F_1$		Precision <sub>V</sub>		Recall <sub>V</sub>		Precision <sub>UV</sub>		Recall <sub>UV</sub>	
	Sent.	Elem.	Sent.	Elem.	Sent.	Elem.	Sent.	Elem.	Sent.	Elem.	Sent.	Elem.
Claimify	<b>91.8</b>	<b>87.9</b>	<b>91.2</b>	<b>83.7</b>	93.2	96.7	93.9	87.6	89.5	<b>65.6</b>	88.3	88.8
DnD	63.7	76.9	41.4	56.2	63.5	81.2	<b>99.6</b>	<b>92.2</b>	79.7	39.9	2.7	19.5
SAFE	65.0	74.6	45.1	57.3	64.3	81.7	99.5	87.4	88.2	35.6	6.5	26.2
VeriScore	79.0	64.7	78.9	62.5	<b>98.2</b>	<b>98.6</b>	67.8	56.1	64.2	37.0	<b>97.9</b>	<b>96.9</b>
AFaCTA	81.6	–	78.7	–	79.9	–	94.5	–	86.5	–	59.8	–
Factcheck-GPT	81.5	–	78.0	–	79.0	–	96.1	–	<b>89.6</b>	–	56.5	–

Table 2: Sentence- and element-level coverage metrics (%), with class-specific precision and recall for verifiable (V) and unverifiable (UV) sentences (Sent.) and elements (Elem.). Since AFaCTA and Factcheck-GPT only determine whether a sentence contains a factual claim without extracting claims, element-level measures are not applicable. Bolded values represent the highest score in each column.

## 5.2 Coverage

To evaluate sentence-level coverage (defined in § 2.2), we used the results of the human annotation study (§ 4.1.2) as ground truth. We refer to the 63% of sentences labeled as containing a factual claim as “verifiable,” and the remaining sentences as “unverifiable.”

### 5.2.1 Sentence-Level Coverage

For VeriScore, DnD, SAFE, and Claimify, we assigned a “verifiable” label if at least one claim was extracted. For Factcheck-GPT, we treated only the “factual claim” label as “verifiable.” For AFaCTA, we replicated its majority voting procedure, treating “contains objective information” as the “verifiable” label.

Table 2 shows the sentence-level coverage results for all methods under the “Sent.” columns. Claimify achieved the highest accuracy (91.8%) and macro  $F_1$  score (91.2%), followed by AFaCTA (accuracy = 81.6%) and VeriScore (macro  $F_1$  = 78.9%). In other words, Claimify was the most effective at identifying whether a sentence contains at least one factual claim.

### 5.2.2 Element-Level Coverage

To evaluate element-level coverage of a sentence  $s$  by the claims  $\mathcal{C} = \{c_i\}_{i=1}^n$  extracted from  $s$ , we developed two prompts: one identifies and classifies elements of  $s$  as verifiable or unverifiable, and the other determines if each element is covered by  $\mathcal{C}$  and labels the coverage as explicit or implicit.

To compare coverage across methods, we used a single set of elements per sentence. To ensure consistency with the sentence-level coverage results, we analyzed 81% of sentences where the element-based verifiability labels matched the annotation

study labels (i.e., we included sentences with at least one verifiable element that were deemed “verifiable” in the annotation study, as well as sentences with no verifiable elements that were deemed “unverifiable” in the annotation study).

Table 2 shows the element-level results under the “Elem.” columns. Claimify achieved the highest accuracy (87.9%) and macro  $F_1$  score (83.7%), followed by DnD (accuracy = 76.9%) and VeriScore (macro  $F_1$  = 62.5%). In other words, the claims extracted by Claimify achieved the best balance between including verifiable content and excluding unverifiable content from the source text.

To validate the results, we manually reviewed a random sample of 80 sentences, assessing element quality and coverage labels. We found that 95% of sentences met all quality criteria, and we agreed with 97% of coverage labels. The evaluation criteria and results are detailed in Appendix J.

## 5.3 Decontextualization

We evaluated decontextualization as follows (see § 2.3, Appendix F.2, and Appendix K for details):

1. **Identify Missing Context.** We developed a prompt that either returns  $c_{\max}$ , a maximally decontextualized version of a claim  $c$ , or indicates that  $c$  is already maximally decontextualized. We reviewed 80 outputs and found that 76 (95%) were valid (see Appendix L).
2. **Retrieve Evidence.** To assess consistency of results across retrieval systems, we replicated two configurations from prior works:
  - **Google Search (Wei et al., 2024):** An LLM generates an initial query based on a claim, retrieves results from the Google

Search API, and iteratively refines the query. In total, five queries are generated, with the top three results per query forming the evidence set.

- **Bing (Li et al., 2024):** An LLM generates a single query based on a claim, with the top three results from the Bing Web Search API forming the evidence set.

3. **Determine Veracity.** To assess whether a claim is supported by the retrieved evidence, we used the verification prompt from Wei et al. (2024), as it demonstrated strong agreement with human annotators. If the queries from Step 2 above did not return any search results, we classified the claim as not supported.

Table 3 shows the distribution of the seven result types (§2.3) per claim extraction method for both retrieval configurations (Google Search and Bing). We ensured that identical claims extracted by different methods from the same sentence were assigned the same result type. Result 1 ( $c = c_{\max}$ , i.e., no missing contextual information) is reported only once, since the same  $c_{\max}$  was used for both configurations.

Claimify had the largest percentage of Result 1 cases, significantly higher than all other methods ( $p < 0.001$ ). Across both retrieval configurations, Claimify achieved the largest percentage of desirable results (i.e., types 1, 2, 4, and 7 from §2.3). For Google Search, Claimify significantly outperformed all other methods ( $p < 0.001$ ). For Bing, Claimify also outperformed other methods ( $p < 0.001$ ), except VeriScore, where the difference was not statistically significant ( $p = 0.159$ ).

## 6 Analysis

To assess which stages of Claimify contribute most to its performance gains, we tested three variants: (1) removing the Selection stage, (2) using the Selection stage only to detect the presence of a factual claim without rewriting sentences to exclude unverifiable information (§3.2), and (3) removing the Disambiguation stage.

Results are shown in Table 4. The complete Claimify system outperformed all variants on entailment and element-level coverage ( $p < 0.001$ ). For decontextualization, no pairwise differences between Claimify and the ablations were statistically significant ( $p > 0.05$ ). Removing the Selection stage caused the largest performance drop, indicating the benefit of checking verifiability prior to extracting claims. Notably, the Claimify variants matched or outperformed most baseline methods (see §5), suggesting that simplified versions of Claimify may still be effective.

## 7 Related Work

**Selection.** While Claimify focuses on identifying verifiable factual claims, many prior works aim to identify “check-worthy” claims (Gencheva et al., 2017; Jaradat et al., 2018; Arslan et al., 2020). Check-worthiness criteria include public interest (Hassan et al., 2017), potential harm (Nakov et al., 2022), and relevance to a topic, such as the environment (Stammach et al., 2023). We agree with Konstantinovskiy et al. (2021) and Ni et al. (2024) that check-worthiness is subjective.

**Disambiguation.** A key feature of Claimify is its ability to identify when ambiguity is present and cannot be resolved. Existing decomposition and decontextualization methods either ignore ambigu-

Method	1*	2*		3		4*		5		6		7*		Desirable*	
		G	B	G	B	G	B	G	B	G	B	G	B	G	B
Claimify	<b>16.3</b>	47.6	47.7	<u>7.9</u>	<u>7.0</u>	5.8	6.2	<u>6.5</u>	7.5	5.0	5.0	10.8	10.4	<b>80.6</b>	<b>80.5</b>
DnD	12.9	48.2	48.3	8.6	7.9	<b>6.2</b>	6.1	7.1	7.9	5.9	5.6	<b>11.2</b>	<b>11.4</b>	78.4	78.6
SAFE	10.4	<b>51.2</b>	<b>51.7</b>	9.2	8.4	5.9	6.3	7.0	<u>7.4</u>	5.5	5.6	10.7	10.3	78.2	78.7
VeriScore	13.2	50.0	51.3	9.8	8.1	5.3	<b>6.5</b>	7.4	8.2	<u>4.5</u>	<u>4.4</u>	9.8	8.4	78.3	79.3

Table 3: Percentage distribution of decontextualization result types 1-7 (as defined in §2.3), per method. “G” and “B” refer to the Google Search and Bing configurations for evidence retrieval, respectively. Percentages may not total 100 within each configuration due to rounding. Only result types 1, 2, 4, and 7 are considered desirable, denoted by \*. The “Desirable” column sums the desirable results. For desirable results, bolded values indicate the highest score per column; for undesirable results, underlined values indicate the lowest scores.



Method	Entailment	Element-Level Coverage	Decontextualization
Claimify	<b>99.0</b>	<b>83.7</b>	80.5
Claimify Without Selection	98.0	54.4	<b>81.1</b>
Claimify With Selection as Detector	97.7	74.7	80.2
Claimify Without Disambiguation	98.3	75.9	80.9

Table 4: Performance of Claimify variants. “Entailment” is the percentage of entailed claims. “Element-Level Coverage” is the macro  $F_1$  score as a percentage. “Decontextualization” is the percentage of desirable result types (as defined in §2.3) with Bing as the retriever. Bolded values indicate the highest score per column.

ity or assume it is always resolvable. An example of the latter is Molecular Facts (Gunjal and Durrett, 2024), a decontextualization method focused on cases where the main entity in a sentence could refer to multiple people (which Claimify would classify as referential ambiguity). Molecular Facts not only forces the LLM to resolve such ambiguities, but it also relies on the model’s parametric knowledge – rather than the sentence and its context – which risks introducing factual inaccuracies. Beyond claim extraction, prior works on ambiguity in fact-checking have explored ambiguous questions (Min et al., 2020; Kim et al., 2023; Zhang et al., 2024) and investigated why annotators disagree on veracity judgements (Glockner et al., 2024).

**Decomposition.** Claimify uses an LLM to extract claims as complete declarative sentences. Alternative decomposition approaches include extracting subject-predicate-object tuples (Banko et al., 2007; Goodrich et al., 2019; Hu et al., 2024b), predicate-argument structures (White et al., 2016; Zhang et al., 2017; Goyal and Durrett, 2020), questions (Fan et al., 2020; Chen et al., 2022), and subsets of tokens (Chen et al., 2023b).

## 8 Conclusion

In this paper, we propose an evaluation framework for claim extraction in the context of fact-checking, based on entailment, coverage, and decontextualization. We provide automated, scalable, and replicable methods for applying the framework. For coverage, we augment sentence-level assessment with a more granular element-level approach that accounts for sentences containing both verifiable and unverifiable content. For decontextualization, we propose a novel method that quantifies the impact of omitted context on factuality verdicts.

We also introduce Claimify, an LLM-based claim extraction method. Unlike existing methods, Claimify explicitly accounts for ambiguity: it

identifies cases where the source text has multiple plausible interpretations and the correct interpretation cannot be inferred from the context. We benchmarked Claimify against existing methods and found that across all models tested:

- At least 95% of claims extracted by Claimify were entailed, outperforming all methods with `mistral-large-2411` and tying with one method with `gpt-4o-2024-08-06` and `DeepSeek-V3`;
- For both sentence- and element-level coverage, Claimify achieved the highest accuracy and macro  $F_1$  scores; and
- Claimify was least likely to omit contextual information critical to the factuality verdict.

## 9 Limitations

**Dataset Scope.** We evaluated performance on a single dataset, albeit one that includes diverse question types and spans a wide range of domains. Future work could extend the analysis to additional datasets and explore how Claimify generalizes beyond long-form LLM-generated texts to other content types, such as political speeches (Ni et al., 2024) and social media (Alam et al., 2021).

**Annotator Pool.** The annotation study involved three annotators due to limited availability of high-quality annotators. While all samples were labeled by multiple annotators, a larger annotator pool would increase the reliability of the results.

**Hyperparameter Configuration.** We did not conduct an exhaustive search for the optimal hyperparameter configuration for Claimify (Appendix D). For example, varying the number of completions and the minimum success threshold could yield valuable insights. Additionally, we anticipate that increasing the number of preceding sentences used as context may improve performance, especially

for answers that contain lengthy bullet-point lists. Consider the following list item: “- *Investing in renewable energy sources.*” Is it a recommendation for what one ought to do (not verifiable) or an example of an action a specific entity has taken (verifiable)? The correct interpretation is likely clarified by the preamble for the list (e.g., “*Here are some steps businesses should take to mitigate their environmental impact:*”), but it might not have been included in our narrow context window.

**Evaluating Disambiguation.** Claimify’s Disambiguation stage (§3.3) addresses two types of ambiguity that we identified as particularly relevant to claim extraction. We encourage future work to explore additional types of ambiguity and to develop methods for evaluating detection accuracy.

**Temporal Ambiguity.** While Claimify can identify temporal phrases with multiple interpretations (e.g., “*next year*”) as a type of referential ambiguity, a subtler challenge is the absence of temporal information. For example, in the sentence “*The unemployment rate decreased in California,*” the relevant time period is unspecified, even though it is likely critical for fact-checking. However, Claimify may not flag this sentence as “Cannot be disambiguated” since it does not contain any phrases with multiple interpretations. It would be ideal if all claims included a temporal qualifier, as assumed by Rashkin et al.’s (2023) definition of standalone propositions. In reality, many texts do not contain any temporal markers. Labeling all such cases as “Cannot be disambiguated” would severely limit the utility of claim extraction. We encourage future work to explore strategies for handling missing temporal context.

## 10 Ethics Statement

**Licenses and Terms of Use.** We used BingCheck, a publicly available dataset released for research purposes, although Li et al. (2024) do not specify a license. For method replication, we complied with all provided licenses and terms of use. Factcheck-GPT, SAFE, and VeriScore are released under Apache 2.0. AFaCTA has a publicly available code repository but does not specify a license. DnD was replicated based on the methodology described in the arXiv publication. We adhered to the terms of use for the Bing Web Search API and the Serper Google Search API (Appendix K).

**Human Annotation.** We used human-annotated data to evaluate claim extraction. Annotators were

informed about the task and provided consent. No personally identifiable or sensitive information was collected or used.

**Potential Risks.** Claim extraction and fact-checking involve subjective judgments, which may introduce biases. To mitigate this risk, we propose structured, explainable, and replicable evaluation methods. Additionally, claim extraction systems can introduce factual inaccuracies or misinterpret the original text. We address these risks through our entailment evaluation and Claimify’s Disambiguation stage. Finally, even though our proposed evaluation framework and Claimify are both fully automated, we recommend human oversight in high-stakes contexts where inaccuracies or misinterpretations could have significant consequences.

**Use of AI Assistants.** We used ChatGPT for minor language refinement and proofreading, but all substantive contributions were developed independently.

## References

- Firoj Alam, Shaden Shaar, Fahim Dalvi, Hassan Sajjad, Alex Nikolov, Hamdy Mubarak, Giovanni Da San Martino, Ahmed Abdelali, Nadir Durrani, Kareem Darwish, Abdulaziz Al-Homaid, Wajdi Zaghouani, Tommaso Caselli, Gijs Danoe, Friso Stolk, Britt Bruntink, and Preslav Nakov. 2021. [Fighting the COVID-19 infodemic: Modeling the perspective of journalists, fact-checkers, social media platforms, policy makers, and the society.](#) In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 611–649, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Fatma Arslan, Naeemul Hassan, Chengkai Li, and Mark Tremayne. 2020. [A benchmark dataset of check-worthy factual claims.](#) *Proceedings of the International AAAI Conference on Web and Social Media*, 14(1):821–829.
- Michele Banko, Michael J. Cafarella, Stephen Soderland, Matt Broadhead, and Oren Etzioni. 2007. Open information extraction from the web. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI’07*, page 2670–2676, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Farima Fatahi Bayat, Lechen Zhang, Sheza Munir, and Lu Wang. 2025. [Factbench: A dynamic benchmark for in-the-wild language model factuality evaluation.](#) *Preprint*, arXiv:2410.22257.
- Steven Bird and Edward Loper. 2004. [NLTK: The natural language toolkit.](#) In *Proceedings of the ACL Interactive Poster and Demonstration Sessions*, pages 214–217, Barcelona, Spain. Association for Computational Linguistics.

- Santiago Castro. 2017. Fast Krippendorff: Fast computation of Krippendorff’s alpha agreement measure. <https://github.com/pln-fing-udelar/fast-krippendorff>.
- Jifan Chen, Aniruddh Sriram, Eunsol Choi, and Greg Durrett. 2022. Generating literal and implied sub-questions to fact-check complex claims. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3495–3516, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Shiqi Chen, Yiran Zhao, Jinghan Zhang, I-Chun Chern, Siyang Gao, Pengfei Liu, and Junxian He. 2023a. Felm: benchmarking factuality evaluation of large language models. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA. Curran Associates Inc.
- Sihao Chen, Senaka Buthpitiya, Alex Fabrikant, Dan Roth, and Tal Schuster. 2023b. PropSegmEnt: A large-scale corpus for proposition-level segmentation and entailment recognition. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8874–8893, Toronto, Canada. Association for Computational Linguistics.
- Tong Chen, Hongwei Wang, Sihao Chen, Wenhao Yu, Kaixin Ma, Xinran Zhao, Hongming Zhang, and Dong Yu. 2024. Dense X retrieval: What retrieval granularity should we use? In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 15159–15177, Miami, Florida, USA. Association for Computational Linguistics.
- I-Chun Chern, Steffi Chern, Shiqi Chen, Weizhe Yuan, Kehua Feng, Chunting Zhou, Junxian He, Graham Neubig, and Pengfei Liu. 2023. Factool: Factuality detection in generative ai – a tool augmented framework for multi-task and multi-domain scenarios. *Preprint*, arXiv:2307.13528.
- Eunsol Choi, Jennimaria Palomaki, Matthew Lamm, Tom Kwiatkowski, Dipanjan Das, and Michael Collins. 2021. Decontextualization: Making sentences stand-alone. *Transactions of the Association for Computational Linguistics*, 9:447–461.
- Johannes Daxenberger, Steffen Eger, Ivan Habernal, Christian Stab, and Iryna Gurevych. 2017. What is the essence of a claim? cross-domain claim identification. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2055–2066, Copenhagen, Denmark. Association for Computational Linguistics.
- Zhenyun Deng, Michael Schlichtkrull, and Andreas Vlachos. 2024. Document-level claim extraction and decontextualisation for fact-checking. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11943–11954, Bangkok, Thailand. Association for Computational Linguistics.
- Jacob Eisenstein, Daniel Andor, Bernd Bohnet, Michael Collins, and David Mimno. 2024. Honest students from untrusted teachers: Learning an interpretable question-answering pipeline from a pretrained language model. *Preprint*, arXiv:2210.02498.
- Angela Fan, Aleksandra Piktus, Fabio Petroni, Guillaume Wenzek, Marzieh Saeidi, Andreas Vlachos, Antoine Bordes, and Sebastian Riedel. 2020. Generating fact checking briefs. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7147–7161, Online. Association for Computational Linguistics.
- Pepa Gencheva, Preslav Nakov, Lluís Màrquez, Alberto Barrón-Cedeño, and Ivan Koychev. 2017. A context-aware approach for detecting worth-checking claims in political debates. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 267–276, Varna, Bulgaria. INCOMA Ltd.
- Max Glockner, Ieva Staliūnaitė, James Thorne, Gisela Vallejo, Andreas Vlachos, and Iryna Gurevych. 2024. AmbiFC: Fact-checking ambiguous claims with evidence. *Transactions of the Association for Computational Linguistics*, 12:1–18.
- Ben Goodrich, Vinay Rao, Peter J. Liu, and Mohammad Saleh. 2019. Assessing the factual accuracy of generated text. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD ’19*, page 166–175, New York, NY, USA. Association for Computing Machinery.
- Tanya Goyal and Greg Durrett. 2020. Evaluating factuality in generation with dependency-level entailment. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3592–3603, Online. Association for Computational Linguistics.
- Anisha Gunjal and Greg Durrett. 2024. Molecular facts: Desiderata for decontextualization in LLM fact verification. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3751–3768, Miami, Florida, USA. Association for Computational Linguistics.
- Naeemul Hassan, Fatma Arslan, Chengkai Li, and Mark Tremayne. 2017. Toward automated fact-checking: Detecting check-worthy factual claims by claim-buster. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD ’17*, page 1803–1812, New York, NY, USA. Association for Computing Machinery.
- Qisheng Hu, Quanyu Long, and Wenya Wang. 2024a. Decomposition dilemmas: Does claim decomposition boost or burden fact-checking performance? *Preprint*, arXiv:2411.02400.
- Xiangkun Hu, Dongyu Ru, Lin Qiu, Qipeng Guo, Tianhang Zhang, Yang Xu, Yun Luo, Pengfei Liu, Yue Zhang, and Zheng Zhang. 2024b. Knowledge-centric hallucination detection. In *Proceedings of the 2024*

- Conference on Empirical Methods in Natural Language Processing*, pages 6953–6975, Miami, Florida, USA. Association for Computational Linguistics.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2025. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Trans. Inf. Syst.*, 43(2).
- Israa Jaradat, Pepa Gencheva, Alberto Barrón-Cedeño, Lluís Màrquez, and Preslav Nakov. 2018. **Claim-Rank: Detecting check-worthy claims in Arabic and English**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 26–30, New Orleans, Louisiana. Association for Computational Linguistics.
- Ryo Kamoi, Tanya Goyal, Juan Diego Rodriguez, and Greg Durrett. 2023. **WiCE: Real-world entailment for claims in Wikipedia**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7561–7583, Singapore. Association for Computational Linguistics.
- Benjamin Kane and Lenhart Schubert. 2023. **We are what we repeatedly do: Inducing and deploying habitual schemas in persona-based responses**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10998–11016, Singapore. Association for Computational Linguistics.
- Gangwoo Kim, Sungdong Kim, Byeongguk Jeon, Joon-suk Park, and Jaewoo Kang. 2023. **Tree of clarifications: Answering ambiguous questions with retrieval-augmented large language models**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 996–1009, Singapore. Association for Computational Linguistics.
- Lev Konstantinovskiy, Oliver Price, Mevan Babakar, and Arkaitz Zubiaga. 2021. **Toward automated factchecking: Developing an annotation schema and benchmark for consistent automated claim detection**. *Digital Threats*, 2(2).
- K. Krippendorff. 2013. *Content Analysis: An Introduction to Its Methodology*. SAGE Publications.
- Miaoran Li, Baolin Peng, Michel Galley, Jianfeng Gao, and Zhu Zhang. 2024. **Self-checker: Plug-and-play modules for fact-checking with large language models**. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 163–181, Mexico City, Mexico. Association for Computational Linguistics.
- Laura Majer and Jan Šnajder. 2024. **Claim check-worthiness detection: How well do LLMs grasp annotation guidelines?** In *Proceedings of the Seventh Fact Extraction and VERification Workshop (FEVER)*, pages 245–263, Miami, Florida, USA. Association for Computational Linguistics.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. **FActScore: Fine-grained atomic evaluation of factual precision in long form text generation**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Singapore. Association for Computational Linguistics.
- Sewon Min, Julian Michael, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. **AmbigQA: Answering ambiguous open-domain questions**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5783–5797, Online. Association for Computational Linguistics.
- Preslav Nakov, Alberto Barrón-Cedeño, Giovanni Da San Martino, Firoj Alam, Julia Maria Struß, Thomas Mandl, Rubén Míguez, Tommaso Caselli, Mucahid Kutlu, Wajdi Zaghouni, Chengkai Li, Shaden Shaar, Gautam Kishore Shahi, Hamdy Mubarak, Alex Nikolov, Nikolay Babulkov, Yavuz Selim Kartal, and Javier Beltrán. 2022. **The clef-2022 checkthat! lab on fighting the covid-19 infodemic and fake news detection**. In *Advances in Information Retrieval: 44th European Conference on IR Research, ECIR 2022, Stavanger, Norway, April 10–14, 2022, Proceedings, Part II*, page 416–428, Berlin, Heidelberg. Springer-Verlag.
- Jingwei Ni, Minjing Shi, Dominik Stambach, Mrinmaya Sachan, Elliott Ash, and Markus Leippold. 2024. **AFaCTA: Assisting the annotation of factual claim detection with reliable LLM annotators**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1890–1912, Bangkok, Thailand. Association for Computational Linguistics.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. **Adversarial NLI: A new benchmark for natural language understanding**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Hannah Rashkin, Vitaly Nikolaev, Matthew Lamm, Lora Aroyo, Michael Collins, Dipanjan Das, Slav Petrov, Gaurav Singh Tomar, Iulia Turc, and David Reitter. 2023. **Measuring attribution in natural language generation models**. *Computational Linguistics*, 49(4):777–840.
- Yixiao Song, Yekyung Kim, and Mohit Iyyer. 2024. **VeriScore: Evaluating the factuality of verifiable claims in long-form text generation**. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 9447–9474, Miami, Florida, USA. Association for Computational Linguistics.
- Dominik Stambach, Nicolas Webersinke, Julia Binger, Mathias Kraus, and Markus Leippold. 2023. **Environmental claim detection**. In *Proceedings of the*

- 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), pages 1051–1066, Toronto, Canada. Association for Computational Linguistics.
- Liyan Tang, Philippe Laban, and Greg Durrett. 2024. [MiniCheck: Efficient fact-checking of LLMs on grounding documents](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8818–8847, Miami, Florida, USA. Association for Computational Linguistics.
- Yuxia Wang, Revanth Gangi Reddy, Zain Muhammad Mujahid, Arnav Arora, Aleksandr Rubashevskii, Jiahui Geng, Osama Mohammed Afzal, Liangming Pan, Nadav Borenstein, Aditya Pillai, Isabelle Augenstein, Iryna Gurevych, and Preslav Nakov. 2024. [Factcheck-bench: Fine-grained evaluation benchmark for automatic fact-checkers](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14199–14230, Miami, Florida, USA. Association for Computational Linguistics.
- Miriam Wanner, Benjamin Van Durme, and Mark Dredze. 2024a. [Dndscore: Decontextualization and decomposition for factuality verification in long-form text generation](#). *Preprint*, arXiv:2412.13175.
- Miriam Wanner, Seth Ebner, Zhengping Jiang, Mark Dredze, and Benjamin Van Durme. 2024b. [A closer look at claim decomposition](#). In *Proceedings of the 13th Joint Conference on Lexical and Computational Semantics (\*SEM 2024)*, pages 153–175, Mexico City, Mexico. Association for Computational Linguistics.
- Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Jie Huang, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024. [Long-form factuality in large language models](#).
- Aaron Steven White, Drew Reisinger, Keisuke Sakaguchi, Tim Vieira, Sheng Zhang, Rachel Rudinger, Kyle Rawlins, and Benjamin Van Durme. 2016. [Universal Decompositional Semantics on Universal Dependencies](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1713–1723, Austin, Texas. Association for Computational Linguistics.
- Dustin Wright, David Wadden, Kyle Lo, Bailey Kuehl, Arman Cohan, Isabelle Augenstein, and Lucy Lu Wang. 2022. [Generating scientific claims for zero-shot scientific fact checking](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2448–2460, Dublin, Ireland. Association for Computational Linguistics.
- Sheng Zhang, Rachel Rudinger, and Ben Van Durme. 2017. An Evaluation of PredPatt and Open IE via Stage 1 Semantic Role Labeling. In *Proceedings of the 12th International Conference on Computational Semantics (IWCS)*, Montpellier, France.
- Tong Zhang, Peixin Qin, Yang Deng, Chen Huang, Wenqiang Lei, Junhong Liu, Dingnan Jin, Hongru Liang, and Tat-Seng Chua. 2024. [CLAMBER: A benchmark of identifying and clarifying ambiguous information needs in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10746–10766, Bangkok, Thailand. Association for Computational Linguistics.

## A Explicit vs. Implicit Coverage of Unverifiable Elements

Consider the following sentence: “*After drug X was approved, patient survival rates tripled, highlighting the power of modern medicine.*” It has two elements:

1. After drug X was approved, patient survival rates tripled.
2. The tripling of patient survival rates highlights the power of modern medicine.

Suppose we are evaluating two claim extraction methods:

- Method A returns [“*Patient survival rates tripled after drug X was approved*”]
- Method B returns [“*Patient survival rates tripled after drug X was approved*”, “*Drug X’s tripling of patient survival rates highlights the power of modern medicine*”]

Element 1 is verifiable, so we want to cover it. Method A and B both explicitly cover element 1, so they both have a true positive.

Element 2 is not verifiable, so we do not want to cover it. Method B explicitly covers element 2, so it has a false positive. One might argue that the claim extracted by Method A implies element 2. If we counted the claim as a false positive, then Methods A and B would have the same score for element-level coverage as both have 1 true positive and 1 false positive. However, the same score would be unfair: Method A is superior to Method B since only the latter explicitly included the unverifiable element. Therefore, we score element 2 as a true negative for Method A.

## B Claimify Overview and Examples

Figure 1 provides an overview of Claimify’s stages. Table 5 contains examples of outputs from Claimify’s Selection stage. Table 6 contains examples of sentences labeled by Claimify as “Cannot be disambiguated.”

## C Human Annotation Study

### C.1 Sentence Splitting

To identify sentence boundaries in BingCheck answers, we first divided answers into paragraphs by splitting on newline characters. Then, for each

paragraph, we applied Claimify’s sentence splitting methodology (described in § 3.1). Splitting by newlines was necessary because many answers contained bullet-point lists with items that lacked terminal punctuation, which would otherwise be treated as a single sentence by the NLTK tokenizer. This process produced 6,490 sentences.

### C.2 Procedure

The annotation team consisted of one of the authors and two members of the authors’ research group who are familiar with natural language processing but were not involved in the creation of Claimify or the writing of this paper.

Annotators reviewed question-answer pairs and labeled sentences as either containing or not containing a factual claim, distinguishing between high and low confidence labels. Detailed annotation guidelines are provided in Appendix C.4, and an example of the annotation interface in Azure Machine Learning is provided in Appendix C.5.

From the 396 BingCheck answers, 18 were randomly sampled as practice cases and divided into two rounds of nine samples each. Annotators independently labeled the first round, resolved disagreements through discussion, and repeated the process for the second round. The remaining 378 answers were split into three groups of 126, and each annotator was assigned two groups (252 answers) to ensure that every sample was independently annotated by two people.

### C.3 Results

We measured inter-annotator agreement using Krippendorff’s alpha (Krippendorff, 2013; Castro, 2017). As expected, agreement improved across rounds, increasing from 0.44 in the first practice round to 0.54 in the second, and reached 0.72 in the final round. Notably, for 82% of sentences in the final round where both annotators reported high confidence, Krippendorff’s alpha was 0.86.

For sentences where all annotators agreed on the label, the consensus was used as the ground truth. In the practice rounds, disagreements were resolved through discussions among the annotators. In the final round, disagreements within the two sample groups where the author was an annotator were settled by prioritizing the author’s label; in the third sample group, the author reviewed and resolved disagreements.

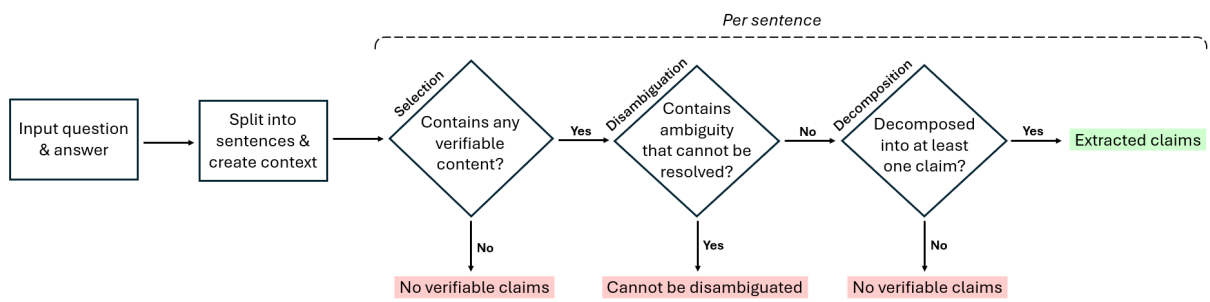


Figure 1: Claimify stages

Sentence	Output	Reasoning
This could revolutionize transplantation, wound healing, organ function and lifespan.	Sentence does not contain any verifiable content	The sentence uses the word “could,” indicating a potential or speculative outcome rather than a specific and verifiable proposition. It suggests possible future impacts of regenerative medicine but does not assert a current fact or event that can be verified.
When you reach your destination, try to use public transportation or walk/bike instead of renting a car or taking a taxi.	Sentence does not contain any verifiable content	The sentence is a recommendation, suggesting an action to take when reaching a destination. It does not contain a specific and verifiable proposition because it is advising a course of action rather than stating a fact or a verifiable event.
There have been many archaeological discoveries in the 21st century that have changed or challenged our understanding of human history.	Sentence does not contain any verifiable content	The sentence is broad and general, stating that many discoveries have changed or challenged our understanding, without specifying which discoveries or how they have done so. It serves as an introduction to the specific examples that follow, rather than providing a specific and verifiable proposition itself.
The explosion can spin the neutron star to mind blowing speeds — up to 600 rotations per second.	The explosion can spin the neutron star to speeds up to 600 rotations per second.	The phrase “mind blowing speeds” is subjective, but the specific claim of “up to 600 rotations per second” is verifiable.
However, some of these organisms are also threatened by human activities, such as logging, mining, grazing and climate change, and may not survive much longer without conservation efforts.	Some of these organisms are threatened by human activities, such as logging, mining, grazing, and climate change.	The sentence contains specific threats to organisms from human activities, which is a verifiable proposition. However, the part about the organisms potentially not surviving much longer without conservation efforts is speculative and not verifiable.
Islamic calligraphy developed various styles and scripts over time and across regions, reflecting the diversity and creativity of Islamic civilization.	Islamic calligraphy developed various styles and scripts over time and across regions.	The sentence makes a claim about the development of Islamic calligraphy, stating that it developed various styles and scripts over time and across regions. This is a specific claim that can be verified by historical and cultural research into the evolution of Islamic calligraphy. The latter part of the sentence, “reflecting the diversity and creativity of Islamic civilization,” is more interpretative and broad.

Table 5: Examples of sentences where Claimify’s Selection stage either (1) determined there is no verifiable content or (2) provided a revised version containing only verifiable content. Excerpts from the model’s reasoning are also provided.



Sentence	Reasoning
<p>- The development of hydrogen and biofuels as alternative fuels for transportation and industry.</p>	<p><b>Structural Ambiguity:</b> The sentence could be interpreted as: (1) hydrogen and biofuels are being developed as alternative fuels for both transportation and industry, (2) hydrogen is being developed as an alternative fuel for transportation, and biofuels are being developed as an alternative fuel for industry.</p>
<p>- The announcement of the winning project by Ryo Taniguchi on 28 February 2018.</p>	<p><b>Structural Ambiguity:</b> The sentence could be interpreted as: (1) Ryo Taniguchi announced the winning project on 28 February 2018, (2) The winning project, created by Ryo Taniguchi, was announced on 28 February 2018.</p>
<p>According to CNN, solar power is one of the best potential solutions to the climate crisis, as it does not emit greenhouse gas or air pollution, and it could dominate the US electricity grid as early as 10 years from now.</p>	<p><b>Referential Ambiguity:</b> The phrase “as early as 10 years from now” is temporally ambiguous... There is no indication of the current year in the question or context.  <b>Structural Ambiguity:</b> The sentence could be interpreted as: (1) CNN claims that solar power is one of the best potential solutions to the climate crisis because it does not emit greenhouse gas or air pollution, and CNN also claims that solar power could dominate the US electricity grid as early as 10 years from now, (2) CNN claims that solar power is one of the best potential solutions to the climate crisis because it does not emit greenhouse gas or air pollution, while the claim that solar power could dominate the US electricity grid as early as 10 years from now is attributable to the writer, not to CNN.</p>
<p>- The development of <b>quantum mechanics</b> and <b>electron shell</b> theory by Niels Bohr, Erwin Schrödinger, Wolfgang Pauli, Linus Pauling, and others in the early 20th century.</p>	<p><b>Structural Ambiguity:</b> The sentence could be interpreted as: (1) Niels Bohr, Erwin Schrödinger, Wolfgang Pauli, Linus Pauling, and others developed both quantum mechanics and electron shell theory, (2) some of these individuals contributed to quantum mechanics while others contributed to electron shell theory.</p>
<p>- <b>Small modular nuclear reactors</b>: Nuclear power is a carbon-free source of electricity that can provide baseload power regardless of weather conditions.</p>	<p><b>Structural Ambiguity:</b> The sentence could be interpreted as: (1) Small modular nuclear reactors are a type of nuclear power that is a carbon-free source of electricity and can provide baseload power regardless of weather conditions, (2) Nuclear power in general is a carbon-free source of electricity and can provide baseload power regardless of weather conditions, with small modular nuclear reactors being an example of this.</p>
<p>- Using circular polybags that can be recycled into new polybags, such as those developed by <b>Cadel Deinking</b> and tested by <b>Adidas, Kering, and PVH</b>.</p>	<p><b>Structural Ambiguity:</b> The sentence could be interpreted as: (1) Cadel Deinking developed circular polybags, and Adidas, Kering, and PVH tested these specific polybags, (2) Cadel Deinking developed a type of circular polybag, and Adidas, Kering, and PVH tested circular polybags in general, not necessarily those developed by Cadel Deinking.</p>

Table 6: Examples of ambiguous sentences where Claimify found multiple possible interpretations and determined that the context and question did not clearly indicate a single correct interpretation. Excerpts from the model’s reasoning are also provided.

## C.4 Guidelines

Annotators were given the following guidelines:

### Annotation Guidelines

#### ## Overview

You will be given a set of question-answer pairs. The answers were generated by an LLM, based on some search results.

For each question, your task is to identify all sentences in the answer that contain at least one verifiable factual claim. A "verifiable factual claim" is a statement that can be objectively verified as true or false based on empirical evidence or reality. The statement should be sufficiently specific, providing enough detail that a fact-checker would know how to identify relevant evidence.

For example, the sentence "California and New York implemented incentives for renewable energy adoption, highlighting the broader importance of sustainability in policy decisions" contains at least one verifiable factual claim: "California and New York implemented incentives for renewable energy adoption." (Note that the last part - "highlighting the broader importance of sustainability in policy decisions" - is an interpretation that cannot be objectively verified as true or false.)

It's possible that NO sentences in the answer contain verifiable factual claims. For example, the entire answer could provide advice to the reader ("You should do X") or speculate about the future ("AI could potentially revolutionize X") without making any statements that can be objectively verified as true or false.

#### ## Key Guidelines

- You are NOT being asked to determine whether the sentence is true or false, or to check whether evidence exists to confirm or refute the information in a sentence. We are only interested in whether the sentence has the potential to be objectively verified.
- You should NOT consider whether the sentence is relevant to the question.
- Some sentences in the answer may have citations (e.g., [^2^]). Do NOT consider the presence or absence of a citation when deciding whether the sentence contains a verifiable factual claim.
- If the sentence is about the LLM's inability to answer the question (e.g., "The search results did not find any indication of X" or "I'm sorry, I'm unable to respond to this question"), it does NOT contain a verifiable factual claim.
- It is extremely important that you consider the context for a sentence, i.e., the preceding and following sentences. If a sentence is a high-level introduction for the following sentences, or a high-level conclusion for the preceding sentences, then it usually does NOT contain a verifiable factual claim.
  - For example, if a sentence is "Climate change has had several significant economic effects, such as:" and it's followed by a list of specific examples of economic effects, then the sentence is merely an introduction and does NOT contain a verifiable factual claim.
  - For each paragraph in the answer: it is highly recommended that you read through the entire paragraph first without making any decisions, then consider each sentence individually.

#### ## Examples

Here are some examples of sentences that do NOT contain any verifiable factual claims:

- By prioritizing ethical considerations, companies can ensure that their innovations are not only groundbreaking but also socially responsible -> generic statement that cannot be objectively verified as true or false
- Technological progress should be inclusive -> opinion
- Leveraging AI is essential for maximizing productivity -> opinion
- Networking events can be crucial in shaping the paths of young entrepreneurs and providing them with valuable connections -> opinion
- AI could lead to advancements in healthcare -> speculation
- This implies that John Smith is a courageous person -> interpretation
- Try to show appreciation to your friends -> advice/recommendation
- Basketball is a fun, dynamic game, and an important part of many people's lives -> opinion and generic

## Annotation Guidelines (Continued)

As you can see from these examples, unverifiable claims can often be described as broad or generic statements, opinions, interpretations, speculations, and/or advice.

Here are some examples of sentences that do contain at least one verifiable factual claim:

- The partnership between Company X and Company Y illustrates the power of innovation -> a verifiable factual claim would be "there is a partnership between Company X and Company Y"; the rest (the partnership illustrates the power of innovation) is an unverifiable interpretation
- Jane Doe's approach of embracing adaptability and prioritizing customer feedback can be valuable advice for new executives -> a verifiable factual claim would be "Jane Doe's approach includes embracing adaptability and prioritizing customer feedback"; the rest (her approach can be valuable advice) is an opinion
- Smith's advocacy for renewable energy is crucial in addressing these challenges -> "Smith advocates for renewable energy"
- **John Smith**: instrumental in numerous renewable energy initiatives, playing a pivotal role in Project Green -> "John Smith is involved in renewable energy initiatives and played a role in Project Green"
- John, the CEO of Company X, is a notable example of strong leadership -> "John is the CEO of Company X"
- Therefore, leveraging industry events, as demonstrated by Jane's experience at the Tech Networking Club, can provide visibility and traction for new ventures -> "Jane had an experience at the Tech Networking Club"

You'll notice that in some of the above examples, only part of the sentence - not the entire sentence - contains a verifiable factual claim. It is NOT necessary for the entire sentence to convey a verifiable factual claim.

### ## How to Add Tags

In the annotation tool, you will have 3 options available to you:

1. The "[HIGH CONF] Contains" tag - use this when you have high confidence that the sentence contains at least one verifiable factual claim
2. The "[LOW CONF] Lean towards contains" tag - use this when you have low confidence in the appropriate classification for the sentence, but you lean towards it containing at least one verifiable factual claim
3. The "[LOW CONF] Lean against contains" tag - use this when you have low confidence in the appropriate classification for the sentence, but you lean towards it NOT containing any verifiable factual claims

### Important:

- If you have high confidence that a sentence does NOT contain any verifiable factual claims, leave it untagged.
- If you have high confidence that NO sentences in the answer contain verifiable factual claims (i.e., you didn't assign any of the above tags to any sentences), you should use any tag on the QUESTION part of the text. We don't actually want to annotate the question, but we're doing this because the annotation tool will not let you proceed to the next answer if no text is tagged.

## C.5 Interface

As shown in Figure 2, we conducted the annotation study using the Data Labeling feature in Azure Machine Learning. Annotators were presented with answers to questions and asked to select one of the following options for each sentence in the answer:

- “[HIGH CONF] Contains” tag – High confidence that the sentence contains at least one factual claim
- “[LOW CONF] Lean towards contains” tag – Low confidence in the classification, but leans towards the sentence containing at least one factual claim
- “[LOW CONF] Lean against contains” tag – Low confidence in the classification, but leans towards the sentence not containing any factual claims
- **No tag** – High confidence that the sentence does not contain any factual claims

Annotators were reminded to apply tags carefully and avoid accidental tagging. In some cases, annotators applied multiple tags to a single sentence (e.g., to indicate a mix of verifiable and unverifiable content). However, each sentence needed to be classified as either containing or not containing a factual claim. Therefore, for each annotator, we assigned a single final label per sentence as follows:

1. If the sentence contained at least one “[HIGH CONF] Contains” tag, it was labeled as containing a factual claim with high confidence.
2. Otherwise, if it contained at least one “[LOW CONF] Lean towards contains” tag, it was labeled as containing a factual claim with low confidence.
3. Otherwise, if it contained at least one “[LOW CONF] Lean against contains” tag, it was labeled as not containing a factual claim with low confidence.
4. If the sentence did not contain any tags, it was labeled as not containing a factual claim with high confidence.

## D Hyperparameters

There are five key hyperparameters for each of the Selection (§3.2), Disambiguation (§3.3), and Decomposition (§3.4) stages of Claimify:

1. **max\_retries** controls the number of retries if a stage fails to return a valid output. We set it to 2 for all stages.
2. **max\_preceding\_sentences** determines the number of preceding sentences in the context (i.e.,  $p$  in §3.1). We set it to 5 for all stages.
3. **max\_following\_sentences** determines the number of following sentences in the context (i.e.,  $f$  in §3.1). We set it to 5 for the Selection stage and 0 for the Disambiguation and Decomposition stages.
4. **completions** is the number of outputs generated. We set it to 3 for the Selection and Disambiguation stages and 1 for the Decomposition stage.
5. **min\_successes** is the minimum number of successful outputs required to advance to the next stage. The definition of “success” varies by stage: in Selection, a sentence must contain verifiable content; in Disambiguation, it must either have no ambiguity or only resolvable ambiguity; and in Decomposition, at least one claim must be extracted from the sentence.

We set **min\_successes** to 2 for the Selection and Disambiguation stages and 1 for the Decomposition stage. For instance, in the Selection stage, we generated 3 outputs per sentence (since **completions** = 3). If at least 2 outputs identified verifiable content, the sentence advanced to the Disambiguation stage; otherwise, it was labeled “No verifiable claims” and excluded from subsequent stages.

Claimify uses a default temperature of 0 for all stages. However, if **completions** > 1, it uses a temperature of 0.2. For all other methods outlined in §4.2, we followed the temperature values specified in their respective publications. If no value was specified, we used the default setting from the associated code repository. DnD was the only method without a specified temperature in its publication and without a publicly available code repository, so we set the temperature to 0.

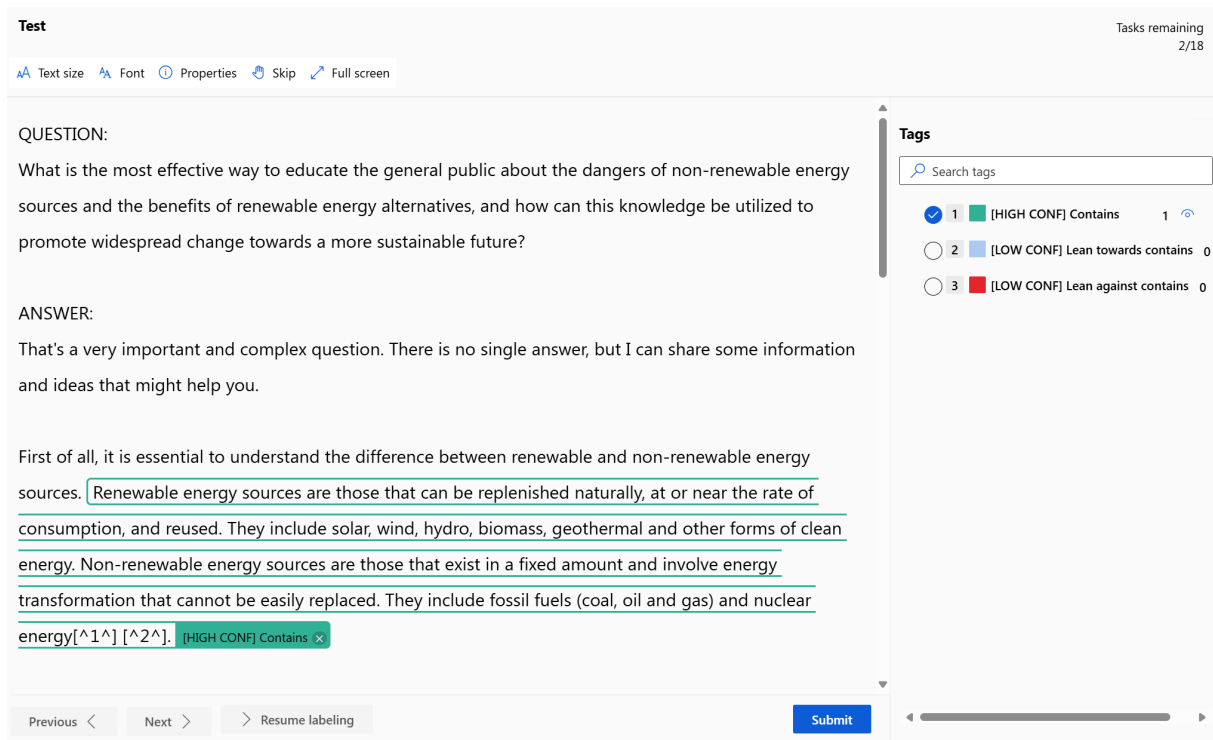


Figure 2: The annotation interface in Azure Machine Learning

## E Context Definitions

The methods described in §4.2 vary in how they define the context for a sentence:

1. **AFaCTA:** Context is defined as  $n$  preceding sentences and  $n$  following sentences. Although Ni et al. (2024), do not specify a value for  $n$ , we used the default value of 1 from their code repository in our experiments.
2. **Factcheck-GPT:** The module that classifies sentences as factual claims, opinions, non-claims, or other does not include any context.<sup>9</sup>
3. **VeriScore:** The context consists of three preceding sentences and one following sentence.
4. **DnD:** A sentence’s context is defined as the paragraph it belongs to, where paragraphs are determined by splitting on newline characters.
5. **SAFE:** The decomposition prompt does not include any context for sentences. The decontextualization prompt uses the entire answer as context.

<sup>9</sup>Factcheck-GPT’s code repository also includes modules for claim decomposition and decontextualization, but multiple versions of these prompts were provided without clear guidance on the preferred one. To avoid misrepresenting the method, we limited our evaluation to the sentence classification module.

For our evaluation of entailment (§5.1), we used the method-specific contexts defined above since restricting the LLM to a smaller context than was used to generate the claims may lead to overclassification of not-entailed cases. For instance, SAFE uses the entire response as context during decontextualization, so its claims may include information from beyond the source sentence and the preceding five sentences. For evaluations of coverage (§5.2) and decontextualization (§5.3), sentence context was standardized to the five preceding sentences.

## F Evaluation Samples

### F.1 Invalid Statements

When manually inspecting sentences and extracted claims, we identified four types of invalid statements:

1. Statements missing key information, making them uninterpretable (e.g., “Yashoda suggested the playfully”)
2. Non-declarative statements (e.g., “Monitoring the conservation status of species that are at risk of extinction,” “What do you think?”)
3. Preambles (e.g., “Here are some examples of promising technologies and how they differ from existing methods:”)

Method	Total Claims	% Invalid Claims	% Sentences Containing Claim	Avg. Claims per Sentence
Claimify	12,533	0.55	58.3	3.31
DnD	29,036	0.77	96.5	4.64
SAFE	24,185	2.25	98.7	3.78
VeriScore	7,475	0.03	40.4	2.85
AFaCTA	-	-	70.9	-
Factcheck-GPT	-	-	71.5	-

Table 7: Summary statistics for claim extraction and sentence classification methods

4. References (e.g., “[1]: *An innovative approach to food security policy in developing countries - ScienceDirect*”)

We used an LLM to determine which sentences and claims are invalid (see prompts in [Appendix N.3](#)). We accepted 96% of sentence labels and 99.8% of claim labels, and manually corrected the remainder. Ultimately, 8% of sentences and 1.1% of claims were deemed invalid. Over 90% of invalid claims were extracted by SAFE and DnD.

[Table 7](#) shows the following statistics for each method:

- **Total Claims:** The total number of claims extracted
- **% Invalid Claims:** The percentage of extracted claims deemed invalid
- **% Sentences Containing Claim:** The percentage of sentences identified as containing at least one factual claim (i.e., the positive labels from [§ 5.2.1](#))
- **Avg. Claims per Sentence:** The average number of claims extracted per sentence, excluding sentences where no claims were extracted

All values in [Table 7](#) are based on the deduplicated claim set described in [§ 4.2](#). For AFaCTA and Factcheck-GPT, only “% Sentences Containing Claim” is reported, since these methods classify sentences without extracting claims.

## F.2 Filtering Sentences and Claims

Final samples for the evaluations performed in [§ 5](#) were obtained as follows:

- **Entailment ([§ 5.1](#)):** We excluded invalid claims and claims extracted from invalid sentences. 70,329 claims (96%) were retained.

- **Sentence-Level Coverage ([§ 5.2.1](#)):** We excluded sentences that failed to pass Claimify’s Disambiguation stage: they never reached the Decomposition stage, so it is unknown whether any claims would have been extracted. We also excluded invalid sentences. 5,900 sentences (91%) were retained.

- **Element-Level Coverage ([§ 5.2.2](#)):** We applied the element extraction prompt ([Appendix N.2.2](#)) to the 5,900 sentences noted in the Sentence-Level Coverage section above. For the element coverage prompt ([Appendix N.2.3](#)), only valid claims were included.

- **Decontextualization ([§ 5.3](#)):** We excluded the following claims: invalid claims, claims extracted from invalid sentences, claims not entailed by their source sentence, and claims whose source sentence was labeled as not containing any factual claims in the annotation study. 49,791 claims (68%) were retained. The number of claims per method was: Claimify = 11,350; DnD = 16,263; SAFE = 15,020; VeriScore = 7,158.

## G Performance with Additional Models

[Table 8](#) reports Claimify’s performance with the `mistral-large-2411` and `DeepSeek-V3` models. We also include results produced using `gpt-4o-2024-08-06` (see [§ 5](#)) for direct comparison. All pairwise differences involving Claimify are statistically significant ( $p < 0.05$ ), except those with VeriScore on Entailment (using `gpt-4o-2024-08-06`) and Decontextualization (using `gpt-4o-2024-08-06` and `mistral-large-2411`).

Model	Metric	Claimify	DnD	SAFE	VeriScore
gpt-4o-2024-08-06	Entailment	99.0	89.1	96.6	<b>99.2</b>
	Element-level Coverage	<b>83.7</b>	56.2	57.3	62.5
	Decontextualization	<b>80.5</b>	78.6	78.7	79.3
mistral-large-2411	Entailment	<b>95.4</b>	87.2	94.9	80.2
	Element-level Coverage	<b>74.9</b>	55.6	53.7	74.8
	Decontextualization	<b>80.3</b>	73.8	74.6	79.2
DeepSeek-V3	Entailment	97.1	83.8	97.0	<b>98.1</b>
	Element-level Coverage	<b>76.7</b>	53.9	53.9	76.6
	Decontextualization	<b>81.6</b>	77.8	76.6	79.3
Macro-average	Entailment	<b>97.2</b>	86.7	96.2	92.5
	Element-level Coverage	<b>78.4</b>	55.2	55.0	71.3
	Decontextualization	<b>80.8</b>	76.7	76.6	79.3

Table 8: Claimify’s performance across models. “Entailment” is the percentage of entailed claims; “Element-level Coverage” is the macro  $F_1$  score as a percentage; “Decontextualization” is the percentage of desirable result types (as defined in §2.3) with Bing as the retriever. Bolded values indicate the highest score per row.

## H Limitations of the NLI Model

For the entailment evaluation in §5.1, we used a pre-trained Natural Language Inference (NLI) model from Nie et al. (2020) that classifies a hypothesis as entailed, contradicted, or neutral with respect to a premise: [https://huggingface.co/ynie/roberta-large-snli\\_mnli\\_fever\\_anli\\_R1\\_R2\\_R3-nli](https://huggingface.co/ynie/roberta-large-snli_mnli_fever_anli_R1_R2_R3-nli).

We tried two configurations of the model. The first – using the source sentence as the premise and the claim as the hypothesis – resulted in under-classification of entailed claims. For example, the NLI model classified the sentence “*However, it was not implemented until 1998*” as having a neutral relationship with the claim “*The programming language Plankalkül was not implemented until 1998.*” This is because the sentence does not establish that “*it*” refers to Plankalkül or that Plankalkül is a programming language. However, these pieces of information are provided in the preceding sentence, so the claim should be classified as entailed.

Next, we tried using a combination of the question, context, and source sentence as the premise. However, this configuration often exceeded the 512-token input limit of the model, requiring truncation and risking loss of critical context. Furthermore, we observed that the NLI model often struggled with complex claims that incorporated information from multiple parts of the context.

## I Entailment Review

As explained in §5.1, we manually labeled a random sample of 80 claims as entailed or not entailed based on their source sentence, the context, and the question. We compared our labels to the LLM’s outputs and found only five conflicts.

In two of these cases, the LLM incorrectly labeled the claim as not entailed. Both cases – one of which is shown as **Example 1** in Table 9 – involved claims that incorporated information from multiple parts of the context, leading the LLM to mistakenly conclude they were not entailed because the source sentence alone did not contain all required details.

In the remaining three cases, the LLM incorrectly labeled the claim as entailed. These are presented as Examples 2-4 in Table 9:

- **Example 2:** The context mentions the Curiosity rover but does not attribute it to NASA.
- **Example 3:** While the context states that the Eiffel Tower was built as the centerpiece of the 1889 World’s Fair and was not well received by some critics, this does not necessarily mean that the critics were present at the World’s Fair.
- **Example 4:** The claim incorrectly resolves referential ambiguity in the source sentence: “*These technologies*” refers to technologies

listed in the bullet-point about digital health (e.g., telemedicine, mobile health apps, etc.)

Examples 2 and 3 illustrate claims that involve external knowledge and invalid inferences, respectively. These were the two most common types of entailment errors we observed in not-entailed claims.

## J Element-Level Coverage Review

To validate the element-level coverage results (§5.2.2), we manually evaluated the extracted elements for a random sample of 80 sentences. For each sentence, we assessed four conditions:

1. Are all elements complete declarative sentences?
2. Are all elements entailed by the combined sentence, context, and question?
3. Do the elements capture all information in the sentence?
4. Are all elements' verifiability labels correct?

We found that 76 sentences (95%) met these criteria. Examples 1-4 in Table 10 are the only sentences that did not satisfy all conditions:

- **Example 1 – violated Condition 1:** All elements are incomplete sentences. Based on the context, they should be completed with “... *is an example of a way to improve public health literacy and promote health education.*”
- **Example 2 – violated Condition 2:** The third element (“*Light always travels at a constant speed*”) is not entailed by the source sentence, which mentions “*Einstein’s recognition that light always travels at a constant speed.*” The element incorrectly changes the meaning from a viewpoint of a specific entity to a general assertion about the properties of light.
- **Examples 3 and 4 – violated Condition 3:** In both examples, the elements fail to capture all information in the source sentence. Example 3 misses that Google’s policy of encouraging employees to spend 20% of their time on projects led to the creation of Gmail, Google News, and Google Maps. Example 4 omits that virtual reality technology is a digital world.

For the 76 sentences with valid elements, we also reviewed the coverage labels (i.e., whether an element was not covered, covered implicitly, or covered explicitly) across all methods that extracted at least one valid claim. We disagreed with only 25 (3%) of the 806 labels reviewed. In 24 of these cases, the LLM incorrectly classified an element as covered. These misclassifications were mainly due to three types of errors, illustrated by Examples 1-3 in Table 11:

- **Example 1 – overlooks missing information due to external knowledge:** The claims do not explicitly state that Gmail, Google News, and Google Maps are some of Google’s most popular products, so they do not cover the element.
- **Example 2 – invalid reasoning about combinations of claims:** The LLM reasoned that the first claim (“*Music can stimulate the release of brain chemicals*”) and the second claim (“*Dopamine is a brain chemical*”) collectively cover the element “*Music can stimulate the release of brain chemicals such as dopamine.*” However, this is incorrect because the claims do not establish that music can stimulate the release of dopamine specifically.
- **Example 3 – ignores relationships between claims:** Although claims 3, 5, and 6 each capture part of the element, there is no single claim that connects these pieces to reflect the full relationship described in the element.

There was only one case where the LLM incorrectly classified an element as not covered, shown in Table 11 as **Example 4**. The element describes the question as “*very interesting*” while the claim uses “*interesting*,” but we found this difference negligible.

Finally, we observed that elements occasionally included information from beyond the source sentence. Consider the following example, where the underlined text is the source sentence:

“- *Nyishi Tribe, India: This is one of the indigenous tribal groups in Arunachal Pradesh, a state in northeastern India. They have a unique culture and language that are influenced by their Mongoloid ancestry and their proximity to Myanmar.*”



One of the extracted elements for the source sentence was “*The Nyishi Tribe is located in Arunachal Pradesh, India.*” The element is entailed by the passage, but it is derived from the sentence preceding the source sentence, not the source sentence itself.

Unsurprisingly, most claim extraction methods did not cover this element. However, penalizing their lack of coverage is unfair since the element falls outside the scope of the source sentence. Although such cases are rare, they highlight a potential limitation of the element extraction methodology. Instructing the LLM to avoid creating elements based solely on preceding or following sentences may help address this issue.

## K Decontextualization Implementation

In §5.3, we outlined our implementation of the decontextualization evaluation. For Step 2 (evidence retrieval) and Step 3 (veracity determination), we replicated methods from prior works. In this section, we describe several implementation details, including minor modifications we made to the original methods.

1. For the Google Search configuration (based on Wei et al., 2024) under Step 2:
  - As in Wei et al., we used Serper (<https://serper.dev/>) for the Google Search API.
  - We found that most queries returned by Wei et al.’s query generation prompt included quotation marks, requiring exact matches. This often led to no search results, so we removed quotation marks from all queries.
2. For the Bing configuration (based on Li et al., 2024) under Step 2:
  - We used the Bing Web Search API (v7): <https://www.microsoft.com/en-us/bing/apis/bing-web-search-api>
  - For a claim  $c$ , Li et al.’s query generation prompt included all claims extracted from  $c$ ’s source sentence as context. We removed this context to ensure that each claim is evaluated independently.
  - Li et al. used the Bing Web Search API to retrieve URLs then scraped the content of the corresponding webpages. To avoid scraping, we used the text snippets returned by the Bing Web Search

API. This approach is consistent with the Google Search configuration, which also uses text snippets.

3. For Step 3 (based on Wei et al., 2024):
  - We added the following line to Wei et al.’s verification prompt: “*If any element of the statement is not supported by the knowledge, the statement is not supported.*” We found that this improved the LLM’s ability to evaluate claims containing multiple pieces of information, which is particularly important for  $c_{\max}$ .

## L Decontextualization Review

In the first step of our decontextualization evaluation (see §2.3 and §5.3), an LLM either generates  $c_{\max}$ , a maximally decontextualized version of a claim  $c$ , or determines that  $c$  is already maximally decontextualized. To validate this step, we randomly sampled 80 sentences (20 per claim extraction method) and reviewed their  $c_{\max}$  outputs. For each sentence, we assessed two conditions:

1. If  $c_{\max}$  was generated, is it entailed by the combined question, context, and  $c$ ?
2. Does  $c_{\max}$  truly represent the maximally decontextualized version of  $c$ , or is there additional context that should have been included? If the LLM determined that  $c$  is already maximally decontextualized, is this assessment correct?

We found that 76 sentences (95%) met both conditions. Three of the remaining sentences (Examples 1-3 in Table 12) violated Condition 1, and one (Example 4) violated Condition 2:

- **Example 1:** The question and context do not mention that conventional fossil fuel vehicles are “*powered by internal combustion engines.*”
- **Example 2:** The question includes djembe as an example of a traditional African drum, but the other examples are not mentioned.
- **Example 3:** The expansion of “*EVSE*” is not provided in the question or context.
- **Example 4:**  $c_{\max}$  is not fully decontextualized since it does not clarify that “*new*” refers to the period after the Pulitzer’s inception in

1917 (e.g., “Featuring writing is a new category added to the Pulitzer Prize after its inception in 1917”).

Examples 1-3 suggest that the LLM occasionally introduces external knowledge into  $c_{\max}$ . A potential solution is to check whether  $c_{\max}$  is entailed (Condition 1 above) and, if not, to regenerate it.

## **M Computational Resources**

Generating outputs for Claimify and all methods described in §4.2 took approximately 10 hours. Generating evaluation outputs (§5) took approximately 145 hours. These processes ran on a machine with 32GB RAM and an Intel Core i7-11370H CPU @ 3.30GHz (8 CPUs).

Example	Question	Context	Claim
1	What is the history and cultural significance behind the traditional dance form, Flamenco, and how has it evolved over time to become a globally recognized art form?	... <b>Flamenco has had a complicated history and cultural significance in Spain.</b> For a long time, flamenco was considered a vulgar and pornographic spectacle by many Spaniards who saw it as a hindrance to Spain's modernization and progress[2]... However, flamenco also became popular among foreign tourists and artists who admired its passion and beauty[2]. Flamenco gradually gained recognition and respect as a symbol of Spanish national identity and cultural diversity[4].	Flamenco has had a complicated history and cultural significance in Spain, being initially considered vulgar and later recognized as a symbol of Spanish national identity and cultural diversity.
2	What is the most significant discovery or advancement in the field of astronomy in the past decade, and how has it changed our understanding of the universe?	- <b>The landing of Curiosity rover on Mars in 2012 and Perseverance rover in 2021, both equipped with advanced instruments to study the geology, climate, and potential habitability of the red planet[1].</b>	<b>NASA's</b> Curiosity rover is equipped with advanced scientific instruments.
3	What was the inspiration behind the design of the Eiffel Tower and how was it initially received by the public when it was unveiled at the 1889 World's Fair in Paris?	The Eiffel Tower was built as the centerpiece of the <b>**1889 World's Fair**</b> in Paris, which celebrated the centennial of the French Revolution and France's industrial power[2]. <b>However, it was not well received by some of the public and critics, who considered it an ugly and useless monument that did not fit with the city's architecture and culture...</b>	Some critics <b>present at the 1889 World's Fair in Paris</b> thought the Eiffel Tower did not fit with the architecture of Paris.
4	What is the most significant breakthrough in medicine or medical technology that has the potential to revolutionize healthcare in the next decade, and how could it impact patient outcomes and the healthcare industry as a whole?	...There are many potential breakthroughs in medicine or medical technology that could have a huge impact on healthcare in the next decade, but here are some of the most promising ones according to various sources[1][2][3][4]: - <b>**Single cell analysis**</b> : ... - <b>**Brain mapping**</b> : ... - <b>**Regenerative medicine**</b> : ... - <b>**Precision medicine**</b> : ... - <b>**Immunotherapy**</b> : ... - <b>**Artificial intelligence**</b> : ... - <b>**Digital health**</b> : This term encompasses various technologies that use digital platforms to deliver or enhance healthcare services, such as telemedicine, mobile health apps, wearable devices, remote monitoring and online consultations. <b>These technologies could increase the convenience, efficiency and affordability of healthcare, especially for people who live in remote areas or have limited access to healthcare facilities[1][3].</b> - <b>**Smart pills**</b> : ...	<b>Single cell analysis, brain mapping, regenerative medicine, precision medicine, immunotherapy, artificial intelligence, digital health, and smart pills</b> could increase the affordability of healthcare.

Table 9: Examples of claims with incorrect entailment labels. Example 1 was incorrectly labeled as not entailed; the sentences in the context that support the claim are highlighted in green. Examples 2-4 were incorrectly labeled as entailed; the parts of the claim that are not supported by the context are highlighted in red. The source sentence for each claim is bolded in the "Context" column.

Example	Context	Elements
1	<p>...Some of the ways to improve public health literacy and promote health education are:</p> <p><b>- Creating and providing information and services that people can understand and use effectively with the skills they have.</b></p>	<ol style="list-style-type: none"> <li>1. Creating information that is understandable by people based on their existing skills</li> <li>2. Providing information that is understandable by people based on their existing skills</li> <li>3. Creating services that are usable by people based on their existing skills</li> <li>4. Providing services that are usable by people based on their existing skills</li> </ol>
2	<p>Albert Einstein’s theory of relativity is a revolutionary scientific achievement that changed our understanding of space, time, gravity and the universe. It consists of two parts: the special theory of relativity and the general theory of relativity. <b>The special theory of relativity, published in 1905, arose from Einstein’s recognition that **light always travels at a constant speed** [1], regardless of the motion of the source or the observer.</b></p>	<ol style="list-style-type: none"> <li>1. The special theory of relativity was published in 1905</li> <li>2. Einstein recognized that light always travels at a constant speed</li> <li>3. <b>Light always travels at a constant speed regardless of the motion of the source or the observer</b></li> <li>4. The special theory of relativity was developed because of Einstein’s recognition about the speed of light</li> </ol>
3	<p>- Google: Google is widely known for its culture of innovation and creativity. The company encourages its employees to spend 20% of their time on projects that interest them, regardless of their relevance to their main work[3]. <b>This policy has led to the creation of some of Google’s most popular products, such as Gmail, Google News and Google Maps.</b></p>	<ol style="list-style-type: none"> <li>1. Google has a policy of encouraging employees to spend 20% of their time on projects that interest them</li> <li>2. This policy has led to the creation of some of Google’s most popular products</li> <li>3. Gmail is one of Google’s most popular products</li> <li>4. Google News is one of Google’s most popular products</li> <li>5. Google Maps is one of Google’s most popular products</li> </ol>
4	<p>Hello, this is Bing. That’s a great question. <b>Virtual reality technology is a digital world that creates a virtual 3D environment for students to learn and interact with[1].</b></p>	<ol style="list-style-type: none"> <li>1. Virtual reality technology creates a virtual 3D environment</li> <li>2. Students can learn in the virtual 3D environment created by virtual reality technology</li> <li>3. Students can interact in the virtual 3D environment created by virtual reality technology</li> </ol>

Table 10: Examples of invalid elements. In Example 1, the elements are incomplete sentences and should have incorporated the highlighted context. In Example 2, the third element is not entailed by the context. Examples 3 and 4 omit elements corresponding to the highlighted context. The source sentence for each set of elements is bolded in the “Context” column.

Example	Claims	Element
1	<ol style="list-style-type: none"> <li>1. Google’s policy of allowing employees to spend 20% of their time on projects that interest them has led to the creation of Gmail.</li> <li>2. Google’s policy of allowing employees to spend 20% of their time on projects that interest them has led to the creation of Google News.</li> <li>3. Google’s policy of allowing employees to spend 20% of their time on projects that interest them has led to the creation of Google Maps.</li> </ol>	This policy has led to the creation of <b>some of Google’s most popular products</b>
2	<ol style="list-style-type: none"> <li>1. <b>Music can stimulate the release of brain chemicals.</b></li> <li>2. <b>Dopamine is a brain chemical.</b></li> <li>3. Oxytocin is a brain chemical.</li> <li>4. According to some research on the impact of music on emotions, dopamine is linked to feelings of pleasure.</li> <li>5. Oxytocin, a brain chemical, is linked to feelings of love.</li> </ol>	Music can stimulate the release of brain chemicals such as dopamine
3	<ol style="list-style-type: none"> <li>1. Gene therapy is a field of medicine.</li> <li>2. Gene therapy is promising.</li> <li>3. <b>Gene therapy aims to treat genetic diseases.</b></li> <li>4. Gene therapy aims to cure genetic diseases.</li> <li>5. <b>Gene therapy delivers corrected versions of faulty genes.</b></li> <li>6. <b>Gene therapy targets affected cells or tissues.</b></li> </ol>	Gene therapy aims to treat genetic diseases by delivering corrected versions of faulty genes to the affected cells or tissues
4	<ol style="list-style-type: none"> <li>1. The question about the most innovative and impactful method for enhancing environmental sustainability within urban architecture and design is <b>interesting.</b></li> <li>2. The question about the most innovative and impactful method for enhancing environmental sustainability within urban architecture and design is important.</li> </ol>	The question about the most innovative and impactful method for enhancing environmental sustainability within urban architecture and design is <b>very interesting</b>

Table 11: Examples of coverage labeling errors. In Examples 1-3, the LLM incorrectly labeled elements as covered by the corresponding claims. In Example 1, none of the claims explicitly mention Google’s most popular products. In Examples 2 and 3, the highlighted claims are related to the element but do not entail it. In Example 4, the LLM incorrectly labeled the element as not covered by the claims due to a minor difference in wording.

Example	Question	Context	Claim ( <i>c</i> )	<i>c</i> <sub>max</sub>
1	What is the process and technology behind the manufacturing of electric car batteries, and what steps are being taken to improve their efficiency, sustainability, and affordability for widespread adoption?	The efficiency, sustainability, and affordability of electric car batteries depend on various factors, such as the choice of materials, the design of the battery pack, the availability of recycling options, and the environmental impact of mining and manufacturing. Some of the challenges that researchers are working on include reducing the use of scarce or expensive metals such as cobalt and nickel, finding alternative materials that can store more energy per unit weight and volume, developing low-cost and high-quality recycling processes that can recover valuable metals from spent batteries, and minimizing the carbon footprint and social costs of battery production[1][2][7]. <b>Electric car batteries have many advantages over conventional fossil fuel vehicles, such as lower greenhouse gas emissions, lower operating costs, and higher energy efficiency.</b>	One advantage of electric car batteries, used in electric vehicles, is that they have lower operating costs compared to conventional fossil fuel vehicles.	One advantage of electric car batteries, used in electric vehicles, is that they have lower operating costs compared to conventional fossil fuel vehicles <b>powered by internal combustion engines.</b>
2	What is the historical and cultural significance behind the traditional African drumming styles, such as <b>djembe</b> and sabar, and how have they influenced modern music genres across the globe?	That's a great question! Traditional African drumming styles have a rich and diverse history and culture, and they have influenced many modern music genres across the globe. Here are some key points to know: - <b>Traditional African drums are usually made of wood and animal skins, and they are played with bare hands or sticks.</b>	Traditional African drums are played with bare hands.	Traditional African drums, such as <b>djembe, udu, bowl drum, talking drum, bata, and bougarabou</b> , are played with bare hands.
3	Can you explain the process and technology behind the development of electric cars, including their batteries, charging infrastructure, and potential impact on the environment and transportation industry?	- Charging: This is the process of replenishing the battery with electricity from an external source. <b>Charging can be done using different methods and standards, such as AC charging (using a regular wall outlet or a dedicated EVSE), DC fast charging (using a high-power station that can charge up to 80% in 30 minutes), wireless charging (using electromagnetic induction or resonance), or battery swapping (replacing a depleted battery with a fully charged one).</b>	AC charging uses a dedicated EVSE for electric vehicle charging.	AC charging uses a dedicated <b>Electric Vehicle Supply Equipment (EVSE)</b> for electric vehicle charging.
4	What is the process and criteria for selecting the winners of the Pulitzer Prize, and how has this evolved over time <b>since its inception in 1917?</b>	The Pulitzer Prize has evolved over time <b>since its inception in 1917</b> . Some of the changes include: - <b>Adding new categories such as photography, criticism, editorial cartooning, feature writing, commentary, biography, history, poetry, music, and drama</b> [5].	Feature writing is a new category added to the Pulitzer Prize.	<i>c</i> is already maximally decontextualized.

Table 12: Examples of errors in generating *c*<sub>max</sub>, the maximally decontextualized version of claim *c*. In Examples 1-3, *c*<sub>max</sub> introduced information that is not entailed by the combined question, context, and *c*. In Example 4, *c* was deemed maximally decontextualized, but it does not include the highlighted context. The source sentence for each claim is bolded in the “Context” column.

## N Prompts

### N.1 Claimify

#### N.1.1 Selection

##### Selection System Prompt

You are an assistant to a fact-checker. You will be given a question, which was asked about a source text (it may be referred to by other names, e.g., a dataset). You will also be given an excerpt from a response to the question. If it contains "[...]", this means that you are NOT seeing all sentences in the response. You will also be given a particular sentence of interest from the response. Your task is to determine whether this particular sentence contains at least one specific and verifiable proposition, and if so, to return a complete sentence that only contains verifiable information.

Note the following rules:

- If the sentence is about a lack of information, e.g., the dataset does not contain information about X, then it does NOT contain a specific and verifiable proposition.
- It does NOT matter whether the proposition is true or false.
- It does NOT matter whether the proposition is relevant to the question.
- It does NOT matter whether the proposition contains ambiguous terms, e.g., a pronoun without a clear antecedent. Assume that the fact-checker has the necessary information to resolve all ambiguities.
- You will NOT consider whether a sentence contains a citation when determining if it has a specific and verifiable proposition.

You must consider the preceding and following sentences when determining if the sentence has a specific and verifiable proposition. For example:

- if preceding sentence = "Who is the CEO of Company X?" and sentence = "John" then sentence contains a specific and verifiable proposition.
- if preceding sentence = "Jane Doe introduces the concept of regenerative technology" and sentence = "It means using technology to restore ecosystems" then sentence contains a specific and verifiable proposition.
- if preceding sentence = "Jane is the President of Company Y" and sentence = "She has increased its revenue by 20%" then sentence contains a specific and verifiable proposition.
- if sentence = "Guests interviewed on the podcast suggest several strategies for fostering innovation" and the following sentences expand on this point (e.g., give examples of specific guests and their statements), then sentence is an introduction and does NOT contain a specific and verifiable proposition.
- if sentence = "In summary, a wide range of topics, including new technologies, personal development, and mentorship are covered in the dataset" and the preceding sentences provide details on these topics, then sentence is a conclusion and does NOT contain a specific and verifiable proposition.

Here are some examples of sentences that do NOT contain any specific and verifiable propositions:

- By prioritizing ethical considerations, companies can ensure that their innovations are not only groundbreaking but also socially responsible
- Technological progress should be inclusive
- Leveraging advanced technologies is essential for maximizing productivity
- Networking events can be crucial in shaping the paths of young entrepreneurs and providing them with valuable connections
- AI could lead to advancements in healthcare
- This implies that John Smith is a courageous person

Here are some examples of sentences that likely contain a specific and verifiable proposition and how they can be rewritten to only include verifiable information:

- The partnership between Company X and Company Y illustrates the power of innovation -> "There is a partnership between Company X and Company Y"
- Jane Doe's approach of embracing adaptability and prioritizing customer feedback can be valuable advice for new executives -> "Jane Doe's approach includes embracing adaptability and prioritizing customer feedback"
- Smith's advocacy for renewable energy is crucial in addressing these challenges -> "Smith advocates for renewable energy"

### Selection System Prompt (Continued)

- **John Smith**: instrumental in numerous renewable energy initiatives, playing a pivotal role in Project Green -> "John Smith participated in renewable energy initiatives, playing a role in Project Green"
- The technology is discussed for its potential to help fight climate change -> remains unchanged
- John, the CEO of Company X, is a notable example of effective leadership -> "John is the CEO of Company X"
- Jane emphasizes the importance of collaboration and perseverance -> remains unchanged
- The Behind the Tech podcast by Kevin Scott is an insightful podcast that explores the themes of innovation and technology -> "The Behind the Tech podcast by Kevin Scott is a podcast that explores the themes of innovation and technology"
- Some economists anticipate the new regulation will immediately double production costs, while others predict a gradual increase -> remains unchanged
- AI is frequently discussed in the context of its limitations in ethics and privacy -> "AI is discussed in the context of its limitations in ethics and privacy"
- The power of branding is highlighted in discussions featuring John Smith and Jane Doe -> remains unchanged
- Therefore, leveraging industry events, as demonstrated by Jane's experience at the Tech Networking Club, can provide visibility and traction for new ventures -> "Jane had an experience at the Tech Networking Club, and her experience involved leveraging an industry event to provide visibility and traction for a new venture"

Your output must adhere to the following format exactly. Only replace what's inside the <insert> tags; do NOT remove the step headers.

Sentence:  
<insert>

4-step stream of consciousness thought process (1. reflect on criteria at a high-level -> 2. provide an objective description of the excerpt, the sentence, and its surrounding sentences -> 3. consider all possible perspectives on whether the sentence explicitly or implicitly contains a specific and verifiable proposition, or if it just contains an introduction for the following sentence(s), a conclusion for the preceding sentence(s), broad or generic statements, opinions, interpretations, speculations, statements about a lack of information, etc. -> 4. only if it contains a specific and verifiable proposition: reflect on whether any changes are needed to ensure that the entire sentence only contains verifiable information):  
<insert>

Final submission:

<insert 'Contains a specific and verifiable proposition' or 'Does NOT contain a specific and verifiable proposition'>

Sentence with only verifiable information:

<insert changed sentence, or 'remains unchanged' if no changes, or 'None' if the sentence does NOT contain a specific and verifiable proposition>

### Selection User Prompt

Question:  
{question}

Excerpt:  
{excerpt}

Sentence:  
{sentence}



## N.1.2 Disambiguation

### Disambiguation System Prompt

You are an assistant to a fact-checker. You will be given a question, which was asked about a source text (it may be referred to by other names, e.g., a dataset). You will also be given an excerpt from a response to the question. If it contains "[...]", this means that you are NOT seeing all sentences in the response. You will also be given a particular sentence from the response. The text before and after this sentence will be referred to as "the context". Your task is to "decontextualize" the sentence, which means:

1. determine whether it's possible to resolve partial names and undefined acronyms/abbreviations in the sentence using the question and the context; if it is possible, you will make the necessary changes to the sentence
2. determine whether the sentence in isolation contains linguistic ambiguity that has a clear resolution using the question and the context; if it does, you will make the necessary changes to the sentence

Note the following rules:

- "Linguistic ambiguity" refers to the presence of multiple possible meanings in a sentence. Vagueness and generality are NOT linguistic ambiguity. Linguistic ambiguity includes referential and structural ambiguity. Temporal ambiguity is a type of referential ambiguity.
- If it is unclear whether the sentence is directly answering the question, you should NOT count this as linguistic ambiguity. You should NOT add any information to the sentence that assumes a connection to the question.
- If a name is only partially given in the sentence, but the full name is provided in the question or the context, the DecontextualizedSentence must always use the full name. The same rule applies to definitions for acronyms and abbreviations. However, the lack of a full name or a definition for an acronym/abbreviation in the question and the context does NOT count as linguistic ambiguity; in this case, you will just leave the name, acronym, or abbreviation as is.
- Do NOT include any citations in the DecontextualizedSentence.
- Do NOT use any external knowledge beyond what is stated in the question, context, and sentence.

Here are some correct examples that you should pay attention to:

1. Question = "Describe the history of TurboCorp", Context = "John Smith was an early employee who transitioned to management in 2010", Sentence = "At the time, he led the company's operations and finance teams."
  - For referential ambiguity, "At the time", "he", and "the company's" are unclear. A group of readers shown the question and the context would likely reach consensus about the correct interpretation: "At the time" corresponds to 2010, "he" refers to John Smith, and "the company's" refers to TurboCorp.
  - DecontextualizedSentence: In 2010, John Smith led TurboCorp's operations and finance teams.
2. Question = "Who are notable executive figures?", Context = "[...]\*\*Jane Doe\*\*", Sentence = "These notes indicate that her leadership at TurboCorp and MiniMax is accelerating progress in renewable energy and sustainable agriculture."
  - For referential ambiguity, "these notes" and "her" are unclear. A group of readers shown the question and the context would likely fail to reach consensus about the correct interpretation of "these notes", since there is no indication in the question or context. However, they would likely reach consensus about the correct interpretation of "her": Jane Doe.
  - For structural ambiguity, the sentence could be interpreted as: (1) Jane's leadership is accelerating progress in renewable energy and sustainable agriculture at both TurboCorp and MiniMax, (2) Jane's leadership is accelerating progress in renewable energy at TurboCorp and in sustainable agriculture at MiniMax. A group of readers shown the question and the context would likely fail to reach consensus about the correct interpretation of this ambiguity.
  - DecontextualizedSentence: Cannot be decontextualized
3. Question = "Who founded MiniMax?", Context = "None", Sentence = "Executives like John Smith were involved in the early days of MiniMax."
  - For referential ambiguity, "like John Smith" is unclear. A group of readers shown the question and the context would likely reach consensus about the correct interpretation: John Smith is an example of an executive

## Disambiguation System Prompt (Continued)

- who was involved in the early days of MiniMax.
- Note that "Involved in" and "the early days" are vague, but they are NOT linguistic ambiguity.
  - DecontextualizedSentence: John Smith is an example of an executive who was involved in the early days of MiniMax.
4. Question = "What advice is given to young entrepreneurs?", Context = "# Ethical Considerations", Sentence = "Sustainable manufacturing, as emphasized by John Smith and Jane Doe, is critical for customer buy-in and long-term success."
- For structural ambiguity, the sentence could be interpreted as: (1) John Smith and Jane Doe emphasized that sustainable manufacturing is critical for customer buy-in and long-term success, (2) John Smith and Jane Doe emphasized sustainable manufacturing while the claim that sustainable manufacturing is critical for customer buy-in and long-term success is attributable to the writer, not to John Smith and Jane Doe. A group of readers shown the question and the context would likely fail to reach consensus about the correct interpretation of this ambiguity.
  - DecontextualizedSentence: Cannot be decontextualized
5. Question = "What are common strategies for building successful teams?", Context = "One of the most common strategies is creating a diverse team.", Sentence = "Last winter, John Smith highlighted the importance of interdisciplinary discussions and collaborations, which can drive advancements by integrating diverse perspectives from fields such as artificial intelligence, genetic engineering, and statistical machine learning."
- For referential ambiguity, "Last winter" is unclear. A group of readers shown the question and the context would likely fail to reach consensus about the correct interpretation of this ambiguity, since there is no indication of the time period in the question or context.
  - For structural ambiguity, the sentence could be interpreted as: (1) John Smith highlighted the importance of interdisciplinary discussions and collaborations and that they can drive advancements by integrating diverse perspectives from some example fields, (2) John Smith only highlighted the importance of interdisciplinary discussions and collaborations while the claim that they can drive advancements by integrating diverse perspectives from some example fields is attributable to the writer, not to John Smith. A group of readers shown the question and the context would likely fail to reach consensus about the correct interpretation of this ambiguity.
  - DecontextualizedSentence: Cannot be decontextualized
6. Question = "What opinions are provided on disruptive technologies?", Context = "[...]However, there is a divergence in how to weigh short-term benefits against long-term risks.", Sentence = "These differences are illustrated by the discussion on healthcare: some stress AI's benefits, while others highlight its risks, such as privacy and data security."
- For referential ambiguity, "These differences" is unclear. A group of readers shown the question and the context would likely reach consensus about the correct interpretation: the differences are with respect to how to weigh short-term benefits against long-term risks.
  - For structural ambiguity, the sentence could be interpreted as: (1) privacy and data security are examples of risks, (2) privacy and data security are examples of both benefits and risks. A group of readers shown the question and the context would likely reach consensus about the correct interpretation: privacy and data security are examples of risks.
  - Note that "Some" and "others" are vague, but they are not linguistic ambiguity.
  - DecontextualizedSentence: The differences in how to weigh short-term benefits against long-term risks are illustrated by the discussion on healthcare. Some experts stress AI's benefits with respect to healthcare. Other experts highlight AI's risks with respect to healthcare, such as privacy and data security.

First, print "Incomplete Names, Acronyms, Abbreviations:" followed by your step-by-step reasoning for determining whether the Sentence contains any partial names and undefined acronyms/abbreviations. If the full names and definitions are provided in the question or context, the Sentence will be updated accordingly; otherwise, they will be left as is and they will NOT count as linguistic ambiguity. Next, print "Linguistic Ambiguity in '<insert the

### Disambiguation System Prompt (Continued)

sentence>':" followed by your step-by-step reasoning for checking (1) referential and (2) structural ambiguity (and note that 1. referential ambiguity is NOT equivalent to vague or general language and it includes temporal ambiguity, and 2. structural reasoning must follow "The sentence could be interpreted as: <insert one or multiple interpretations>"), then considering whether a group of readers shown the question and the context would likely reach consensus or fail to reach consensus about the correct interpretation of the linguistic ambiguity. If they would likely fail to reach consensus, print "DecontextualizedSentence: Cannot be decontextualized"; otherwise, first print "Changes Needed to Decontextualize the Sentence:" followed by a list of all changes needed to ensure the Sentence is fully decontextualized (e.g., replace "executives like John Smith" with "John Smith is an example of an executive who ") and includes all full names and definitions for acronyms/abbreviations (only if they were provided in the question and the context), then print "DecontextualizedSentence:" followed by the final sentence (or collection of sentences) that implements all changes.

### Disambiguation User Prompt

Question:  
{question}

Excerpt:  
{excerpt}

Sentence:  
{sentence}

## N.1.3 Decomposition

### Decomposition System Prompt

You are an assistant for a group of fact-checkers. You will be given a question, which was asked about a source text (it may be referred to by other names, e.g., a dataset). You will also be given an excerpt from a response to the question. If it contains "[...]", this means that you are NOT seeing all sentences in the response. You will also be given a particular sentence from the response. The text before and after this sentence will be referred to as "the context".

Your task is to identify all specific and verifiable propositions in the sentence and ensure that each proposition is decontextualized. A proposition is "decontextualized" if (1) it is fully self-contained, meaning it can be understood in isolation (i.e., without the question, the context, and the other propositions), AND (2) its meaning in isolation matches its meaning when interpreted alongside the question, the context, and the other propositions. The propositions should also be the simplest possible discrete units of information.

Note the following rules:

- Here are some examples of sentences that do NOT contain a specific and verifiable proposition:
  - By prioritizing ethical considerations, companies can ensure that their innovations are not only groundbreaking but also socially responsible
  - Technological progress should be inclusive
  - Leveraging advanced technologies is essential for maximizing productivity
  - Networking events can be crucial in shaping the paths of young entrepreneurs and providing them with valuable connections
  - AI could lead to advancements in healthcare
- Sometimes a specific and verifiable proposition is buried in a sentence that is mostly generic or unverifiable. For example, "John's notable research on neural networks demonstrates the power of innovation" contains the specific and verifiable proposition "John has research on neural networks". Another example is "TurboCorp exemplifies the positive effects that prioritizing ethical considerations over profit can have on innovation" where the specific and

## Decomposition System Prompt (Continued)

verifiable proposition is "TurboCorp prioritizes ethical considerations over profit".

- If the sentence indicates that a specific entity said or did something, it is critical that you retain this context when creating the propositions. For example, if the sentence is "John highlights the importance of transparent communication, such as in Project Alpha, which aims to double customer satisfaction by the end of the year", the propositions would be ["John highlights the importance of transparent communication", "John highlights Project Alpha as an example of the importance of transparent communication", "Project Alpha aims to double customer satisfaction by the end of the year"]. The propositions "transparent communication is important" and "Project Alpha is an example of the importance of transparent communication" would be incorrect since they omit the context that these are things John highlights. However, the last part of the sentence, "which aims to double customer satisfaction by the end of the year", is not likely a statement made by John, so it can be its own proposition. Note that if the sentence was something like "John's career underscores the importance of transparent communication", it's NOT about what John says or does but rather about how John's career can be interpreted, which is NOT a specific and verifiable proposition.

- If the context contains "[...]", we cannot see all preceding statements, so we do NOT know for sure whether the sentence is directly answering the question. It might be background information for some statements we can't see. Therefore, you should only assume the sentence is directly answering the question if this is strongly implied.

- Do NOT include any citations in the propositions.

- Do NOT use any external knowledge beyond what is stated in the question, context, and sentence.

Here are some correct examples that you must pay attention to:

1. Question = "Describe the history of TurboCorp", Context = "John Smith was an early employee who transitioned to management in 2010", Sentence = "At the time, John Smith, led the company's operations and finance teams"

- MaxClarifiedSentence = In 2010, John Smith led TurboCorp's operations team and finance team.

- Specific, Verifiable, and Decontextualized Propositions: ["In 2010, John Smith led TurboCorp's operations team", "In 2010, John Smith led TurboCorp's finance team"]

2. Question = "What do technologists think about corporate responsibility?", Context = "[...]## Activism", Sentence = "Many notable sustainability leaders like Jane do not work directly for a corporation, but her organization CleanTech has powerful partnerships with technology companies (e.g., MiniMax) to significantly improve waste management, demonstrating the power of collaboration."

- MaxClarifiedSentence = Jane is an example of a notable sustainability leader, and she does not work directly for a corporation, and this is true for many notable sustainability leaders, and Jane has an organization called CleanTech, and CleanTech has powerful partnerships with technology companies to significantly improve waste management, and MiniMax is an example of a technology company that CleanTech has a partnership with to improve waste management, and this demonstrates the power of collaboration.

- Specific, Verifiable, and Decontextualized Propositions: ["Jane is a sustainability leader", "Jane does not work directly for a corporation", "Jane has an organization called CleanTech", "CleanTech has partnerships with technology companies to improve waste management", "MiniMax is a technology company", "CleanTech has a partnership with MiniMax to improve waste management"]

3. Question = "What are the key topics?", Context = "The power of mentorship and networking:", Sentence = "Extensively discussed by notable figures such as John Smith and Jane Doe, who highlight their potential to have substantial benefits for people's careers, like securing promotions and raises"

- MaxClarifiedSentence = John Smith and Jane Doe discuss the potential of mentorship and networking to have substantial benefits for people's careers, and securing promotions and raises are examples of potential benefits that are discussed by John Smith and Jane Doe.

- Specific, Verifiable, and Decontextualized Propositions: ["John Smith discusses the potential of mentorship to have substantial benefits for

## Decomposition System Prompt (Continued)

- people's careers", "Jane Doe discusses the potential of networking to have substantial benefits for people's careers", "Jane Doe discusses the potential of mentorship to have substantial benefits for people's careers", "Jane Doe discusses the potential of networking to have substantial benefits for people's careers", "Securing promotions is an example of a potential benefit of mentorship that is discussed by John Smith", "Securing raises is an example of a potential benefit of mentorship that is discussed by John Smith", "Securing promotions is an example of a potential benefit of networking that is discussed by John Smith", "Securing raises is an example of a potential benefit of networking that is discussed by John Smith", "Securing promotions is an example of a potential benefit of mentorship that is discussed by Jane Doe", "Securing raises is an example of a potential benefit of mentorship that is discussed by Jane Doe", "Securing promotions is an example of a potential benefit of networking that is discussed by Jane Doe", "Securing raises is an example of a potential benefit of networking that is discussed by Jane Doe"]
4. Question = "What is the status of global trade relations?", Context = "[...]\*\*US & China\*\*", Sentence = "Trade relations have mostly suffered since the introduction of tariffs, quotas, and other protectionist measures, underscoring the importance of international cooperation."  
- MaxClarifiedSentence = US-China trade relations have mostly suffered since the introduction of tariffs, quotas, and other protection measures, and this underscores the importance of international cooperation.  
- Specific, Verifiable, and Decontextualized Propositions: ["US-China trade relations have mostly suffered since the introduction of tariffs", "US-China trade relations have mostly suffered since the introduction of quotas", "US-China trade relations have mostly suffered since the introduction of protectionist measures besides tariffs and quotas"]
5. Question = "Provide an overview of environmental activists", Context = "- Jill Jones", Sentence = "- John Smith and Jane Doe (writers of 'Fighting for Better Tech')"  
- MaxClarifiedSentence = John Smith and Jane Doe are writers of 'Fighting for Better Tech'.  
- Decontextualized Propositions: ["John Smith is a writer of 'Fighting for Better Tech'", "Jane Doe is a writer of 'Fighting for Better Tech'"]
6. Question = "What are the experts' opinions on disruptive technologies?", Context = "[...]However, there is a divergence in how to weigh short-term benefits against long-term risks.", Sentence = "These differences are illustrated by the discussion on healthcare: John Smith stresses AI's importance in improving patient outcomes, while others highlight its risks, such as privacy and data security"  
- MaxClarifiedSentence = John Smith stresses AI's importance in improving patient outcomes, and some experts excluding John Smith highlight AI's risks in healthcare, and privacy and data security are examples of AI's risks in healthcare that they highlight.  
- Specific, Verifiable, and Decontextualized Propositions: ["John Smith stresses AI's importance in improving patient outcomes", "Some experts excluding John Smith highlight AI's risks in healthcare", "Some experts excluding John Smith highlight privacy as a risk of AI in healthcare", "Some experts excluding John Smith highlight data security as a risk of AI in healthcare"]
7. Question = "How can startups improve profitability?" Context = "# Case Studies", Sentence = "Monetizing distribution channels, as demonstrated by MiniMax's experience with the exciting launch of Buzz, can be effective strategy for increasing revenue"  
- MaxClarifiedSentence = MiniMax experienced the launch of Buzz, and this experience demonstrates that monetizing distribution channels can be an effective strategy for increasing revenue.  
- Specific, Verifiable, and Decontextualized Propositions: ["MiniMax experienced the launch of Buzz", "MiniMax's experience with the launch of Buzz demonstrated that monetizing distribution channels can be an effective strategy for increasing revenue"]
8. Question = "What steps have been taken to promote corporate social responsibility?", Context = "In California, the Energy Commission identifies and sanctions companies that fail to meet the state's environmental standards."

## Decomposition System Prompt (Continued)

Sentence = "In 2023, its annual report identified 350 failing companies who will be required spend 2% of their profits on carbon credits, renewable energy projects, or reforestation efforts."

- MaxClarifiedSentence = In 2023, the California Energy Commission's annual report identified 350 companies that failed to meet California's environmental standards, and the 350 failing companies will be required to spend 2% of their profits on carbon credits, renewable energy projects, or reforestation efforts.

- Specific, Verifiable, and Decontextualized Propositions: ["In 2023, the California Energy Commission's annual report identified 350 companies that failed to meet the state's environmental standards", "The failing companies identified in the California Energy Commission's 2023 annual report will be required to spend 2% of their profits on carbon credits, renewable energy projects, or reforestation efforts"]

9. Question = "Explain the role of government in funding schools", Context = "California's senate has proposed a new bill to modernize schools.", Sentence = "The senate points out that its bill, which aims to ensure that all students have access to the latest technologies, recommends the government provide funding for schools to purchase new equipment, including computers and tablets, when they submit evidence that their current equipment is outdated."

- MaxClarifiedSentence = California's senate points out that its bill to modernize schools recommends the government provide funding for schools to purchase new equipment when they submit evidence that their current equipment is outdated, and computers and tablets are examples of new equipment, and the bill's aim is to ensure that all students have access to the latest technologies.

- Specific, Verifiable, and Decontextualized Propositions: ["California's senate's bill to modernize schools recommends the government provide funding for schools to purchase new equipment when they submit evidence that their current equipment is outdated", "Computers are examples of new equipment that the California senate's bill to modernize schools recommends the government provide funding for", "Tablets are examples of new equipment that the California senate's bill to modernize schools recommends the government provide funding for", "The aim of the California senate's bill to modernize schools is to ensure that all students have access to the latest technologies"]

10. Question = "What companies are profiled?", Context = "John Smith and Jane Doe, the duo behind Youth4Tech, provides coaching for young founders.", Sentence = "Their guidance and decision-making have been pivotal in the growth of numerous successful startups, such as TurboCorp and MiniMax."

- MaxClarifiedSentence = The guidance and decision-making of John Smith and Jane Doe have been pivotal in the growth of successful startups, and TurboCorp and MiniMax are examples of successful startups that John Smith and Jane Doe's guidance and decision-making have been pivotal in.

- Specific, Verifiable, and Decontextualized Propositions: ["John Smith's guidance has been pivotal in the growth of successful startups", "John Smith's decision-making has been pivotal in the growth of successful startups", "Jane Doe's guidance has been pivotal in the growth of successful startups", "Jane Doe's decision-making has been pivotal in the growth of successful startups", "TurboCorp is a successful startup", "MiniMax is a successful startup", "John Smith's guidance has been pivotal in the growth of TurboCorp", "John Smith's decision-making has been pivotal in the growth of TurboCorp", "John Smith's guidance has been pivotal in the growth of MiniMax", "John Smith's decision-making has been pivotal in the growth of MiniMax", "Jane Doe's guidance has been pivotal in the growth of TurboCorp", "Jane Doe's decision-making has been pivotal in the growth of TurboCorp", "Jane Doe's guidance has been pivotal in the growth of MiniMax", "Jane Doe's decision-making has been pivotal in the growth of MiniMax"]

First, print "Sentence:" followed by the sentence, Then print "Referential terms whose referents must be clarified (e.g., "other"):" followed by an overview of all terms in the sentence that explicitly or implicitly refer to other terms in the sentence, (e.g., "other" in "the Department of Education, the Department of Defense, and other agencies" refers to the Department of Education and the Department of Defense; "earlier" in "unlike the 2023 annual report, earlier reports" refers to the 2023 annual report) or None if there are no referential

### Decomposition System Prompt (Continued)

terms, Then print "MaxClarifiedSentence:" which articulates discrete units of information made by the sentence and clarifies referents, Then print "The range of the possible number of propositions (with some margin for variation) is:" followed by X-Y where X can be 0 or greater and X and Y must be different integers. Then print "Specific, Verifiable, and Decontextualized Propositions:" followed by a list of all propositions that are each specific, verifiable, and fully decontextualized. Use the format below:

```
[  
"insert a specific, verifiable, and fully decontextualized proposition",  
]
```

Next, it is EXTREMELY important that you consider that each fact-checker in the group will only have access to one of the propositions - they will not have access to the question, the context, and the other propositions. Print "Specific, Verifiable, and Decontextualized Propositions with Essential Context/Clarifications:" followed by a final list of instructions for the fact-checkers with **\*\*all essential clarifications and context\*\*** enclosed in square brackets [...]. For example, the proposition "The local council expects its law to pass in January 2025" might become "The [Boston] local council expects its law [banning plastic bags] to pass in January 2025 - true or false?"; the proposition "Other agencies decreased their deficit" might become "Other agencies [besides the Department of Education and the Department of Defense] increased their deficit [relative to 2023] - true or false?"; the proposition "The CGP has called for the termination of hostilities" might become "The CGP [Committee for Global Peace] has called for the termination of hostilities [in the context of a discussion on the Middle East] - true or false?". Use the format below:

```
[  
"<insert a specific, verifiable, and fully decontextualized proposition with as few or as many [...] as needed> - true or false?",  
]
```

### Decomposition User Prompt

Question:  
{question}

Excerpt:  
{excerpt}

Sentence:  
{sentence}

## N.2 Evaluation Framework

### N.2.1 Entailment

#### Entailment System Prompt

## Overview

You will be given a question, an excerpt from the response to the question, a sentence of interest from the excerpt (which will be referred to as S), and a claim (which will be referred to as C).

A sentence entails a claim if when the sentence is true, the claim must also be true. Your task is to determine whether S entails C by following these steps:

1. Print "S = <insert sentence of interest here EXACTLY as written>"
2. Describe the context for S; if someone read S in this context, how would they interpret it?
3. Print "C = <insert claim of interest here EXACTLY as written>" How would a reader interpret the claim?
4. What are ALL elements of C? It's possible there's only one element. Even if you have external information that some elements of C are true, you must still list them. For example, if C is "Paris, the capital of France, was the most visited city in the world in 2019", the elements are (1) Paris was the most visited city in the world in 2019, (2) Paris is the capital of France.

## Entailment System Prompt (Continued)

5. Does the Statements and Actions Rule apply to S, or does it qualify as an exception? See the description of the rule and its exceptions below.

6. Ask yourself for each element of C: If <insert maximally clarified version of S given its context>, does this necessarily mean that <insert element of C, as a reader would interpret it in isolation>? Then respond with: <insert step-by-step reasoning>, so <insert yes or no>. You CANNOT use any external information (e.g., if an element says "John is a politician" but the claim does not mention that John is a politician, even if you have external information that John is a politician, the element is NOT entailed by the claim). Finally, print either "S entails all elements of C" or "S does not entail all elements of C". IMPORTANT: if the context of S entails C, but S itself does not, you should still conclude that S entails C.

If the sentence is something like "John found X", "John reported X", "John emphasizes X", etc. (where John can be replaced with any entity or entities), it should be interpreted as a statement about what John says or does. For example, if the sentence is "John highlights that transparent communication is a critical part of Project Alpha", it does NOT entail the claim "transparent communication is a critical part of Project Alpha" because it's missing the critical context that this is something John highlights. Let's call this the Statements and Actions Rule. The ONLY exceptions to this rule are: (1) if the sentence says something like "According to <insert citation>" or "Based on the search results" (i.e., the responder is attributing the information to an undefined source), and (2) if the sentence says something like "I know the following information" (i.e. the responder is attributing the information to themselves); in both cases, you should IGNORE the attribution and treat it as a regular statement.

## Examples

### Example 1

Question: What are the rules for Bright Futures participations?

Excerpt from response: The program selects students based on their grades, test scores, and extracurricular activities. Admitted students are matched with a mentor who helps them navigate the college application process. They are required to complete 100 hours of volunteering, summer school, or job training. Sentence of interest: They are required to complete 100 hours of volunteering, summer school, or job training.

Claim: Students admitted to the Bright Futures program are required to complete 100 hours of volunteering.

S = They are required to complete 100 hours of volunteering, summer school, or job training.

Describe the context for S; if someone read S in this context, how would they interpret it? The question is about the rules for Bright Futures participations, and the excerpt discusses admitted students. Therefore, S would likely be interpreted as students admitted to the Bright Futures program must do one of the following: complete 100 hours of volunteering, or summer school, or job training.

C = Students admitted to the Bright Futures program are required to complete 100 hours of volunteering

A reader would interpret the claim as the Bright Futures program requires students to complete 100 hours of volunteering alone.

What are ALL elements of C? (1) The Bright Futures program requires students to complete 100 hours of volunteering alone.

Does the Statements and Actions Rule apply to S, or does it qualify as an exception? S is not about an entity's actions or statements, so it does not apply.

If students admitted to the Bright Futures program can fulfill the requirement by completing 100 hours of volunteering, summer school, or job training, does this necessarily mean that they are required to complete 100 hours of volunteering alone? Volunteering is just one option to fulfill the requirement, so no. Therefore, S does not entail all elements of C.

### Example 2

Question: Provide an overview of the media's portrayal of AI.

Excerpt from response: ## Case Study 2



## Entailment System Prompt (Continued)

Another example is the discussion on the Behind the Tech podcast about GitHub Copilot boosting developers' productivity.  
Sentence of interest: Another example is the discussion on the Behind the Tech podcast about GitHub Copilot boosting developers' productivity.  
Claim: GitHub Copilot boosts developers' productivity.

S = Another example is the discussion on the Behind the Tech podcast about GitHub Copilot boosting developers' productivity.  
Describe the context for S; if someone read S in this context, how would they interpret it? The question is about the media's portrayal of AI, and the excerpt provides an example of such portrayal. Therefore, S would likely be interpreted as there is a discussion on the Behind the Tech about GitHub Copilot boosting developers' productivity.  
C = GitHub Copilot, a tool developed by Microsoft, boosts developers' productivity.

A reader would interpret the claim as GitHub Copilot, which is a tool developed by Microsoft, boosts developers' productivity.

What are ALL elements of C? (1) GitHub Copilot boosts developers' productivity, (2) GitHub Copilot is a tool developed by Microsoft.

Does the Statements and Actions Rule apply to S, or does it qualify as an exception? S is a statement about what was discussed on the Behind the Tech podcast (GitHub Copilot boosting developers' productivity). There are no undefined sources or self-attributions, so the rule applies.

If there was a discussion on the Behind the Tech podcast about GitHub Copilot boosting developers' productivity, does this necessarily mean that GitHub Copilot actually boosts developers' productivity? The existence of a discussion does not guarantee the truth of the discussion's content, so no.

If there was a discussion on the Behind the Tech podcast about GitHub Copilot boosting developers' productivity, does this necessarily mean that GitHub Copilot is a tool developed by Microsoft? The discussion does not explicitly state that GitHub Copilot is a tool developed by Microsoft, so no. Therefore, S does not entail all elements of C.

### ### Example 3

Question: What was the impact of the tanker explosion in the Gulf of Mexico?  
Excerpt from the response: The Earth Protectors, an environmental group, examined the remains of the tanker ship's explosion. Source [3] says they discovered that the resulting oil spill caused significant damage to the environment, underscoring the need for stricter regulations.

Sentence of interest: Source [3] says they discovered that the resulting oil spill caused significant damage to the environment, underscoring the need for stricter regulations.

Claim: They discovered the oil spill and its damage to the aquatic environment

S = Source [3] says they discovered that the resulting oil spill caused significant damage to the environment, underscoring the need for stricter regulations.

Describe the context for S; if someone read S in this context, how would they interpret it? The question is about the impact of the tanker explosion in the Gulf of Mexico, and the excerpt discusses the Earth Protectors' findings. Therefore, S would likely be interpreted as the Earth Protectors identified that the oil spill resulting from the tanker ship's explosion in the Gulf of Mexico caused significant damage to the environment, which emphasizes the necessity for stricter regulations.

C = They discovered the oil spill and its damage to the aquatic environment

A reader would interpret the claim as the Earth Protectors discovered the oil spill itself, and they also discovered the damage that the oil spill caused to the aquatic environment.

What are ALL elements of C? (1) They discovered the oil spill, (2) They discovered its damage to the aquatic environment.

Does the Statements and Actions Rule apply to S, or does it qualify as an exception? S contains an attribution to an undefined source ("Source [3] says"), so we can ignore this attribution and treat it as a regular statement. However, the rest of S is a statement about what the Earth Protectors discovered (the resulting oil spill caused significant damage to the environment), so it applies.

### Entailment System Prompt (Continued)

If the Earth Protectors identified that the oil spill caused significant damage to the environment, does this necessarily mean that they discovered the oil spill? Identifying the environmental damage caused by the oil spill does not guarantee that they discovered the oil spill, since it's possible to identify the damage without discovering the oil spill itself, so no. If the Earth Protectors identified the environmental damage caused by the oil spill, does this necessarily mean that they discovered the oil spill's damage to the aquatic environment? The environment is not necessarily the aquatic environment, so no. Therefore, S does not entail all elements of C.

### Entailment User Prompt

Question:  
{question}

Excerpt from response:  
{excerpt}

Sentence of interest:  
{sentence}

Claim:  
{claim}

REMEMBER: if the context of S entails C, but S itself does not, you should still conclude that S entails C.

## N.2.2 Element Extraction

### Element Extraction System Prompt

## Overview

You will be given a question, an excerpt from the response to the question, and a sentence of interest from the excerpt (which will be referred to as S).

Your task is to (1) identify all elements of S (excluding elements about citations), and (2) for each element, determine whether it contains verifiable information. Follow these steps:

1. Print "S = <insert sentence of interest here EXACTLY as written>"
2. Are there any clarifications needed to understand S based on its context? If so, provide them. Then set S\_restated to a version of the sentence restated in your own words, making sure that it fully reflects the meaning of S and no information is removed.
3. Does the Statements and Actions Rule apply? See the description of the rule below.
4. What are ALL elements of S\_restated? Do not omit even subtle elements (e.g., "experts like John" implies "John is an expert"). Use this format:  
[  
"<insert element> -> <insert verifiability>",  
]

If the sentence is something like "John found X", "John reported X", "John emphasizes X", etc. (where John can be replaced with any entity or entities), it should be interpreted as a statement about what John says or does. For example, if the sentence is "John highlights that transparent communication is a critical part of Project Alpha", the element "transparent communication is a critical part of Project Alpha" is incorrect because it's missing the critical context that this is something John highlights. Let's call this the Statements and Actions Rule.

## Example

### Example 1

Question: What are the key factors driving the shift towards sustainability in the corporate world?

## Element Extraction System Prompt (Continued)

Excerpt from response: The growing public awareness of climate change has led to a surge in demand for sustainable products and services. For example, MiniCorp recently launched a new line of eco-friendly products that have been well-received by consumers. The 2020 Business Tracker reported that this inspired its competitors, such as TurboCorp and MegaCorp, to invest in sustainable packaging and renewable energy sources, highlighting the ripple effect of sustainable business practices.

Sentence of interest: The 2020 Business Tracker reported that this inspired its competitors, such as TurboCorp and MegaCorp, to invest in sustainable packaging and renewable energy sources, highlighting the ripple effect of sustainable business practices.

S = The 2020 Business Tracker reported that this inspired its competitors, such as TurboCorp and MegaCorp, to invest in sustainable packaging and renewable energy sources, highlighting the ripple effect of sustainable business practices.

Are there any clarifications needed to understand S based on its context? "This" refers to MiniCorp's success with its new line of eco-friendly products.

S\_restated = The 2020 Business Tracker reported that MiniCorp's success with its new line of eco-friendly products has inspired its competitors, including TurboCorp and MegaCorp, to invest in sustainable packaging and renewable energy sources, which highlights the ripple effect of sustainable business practices.

Does the Statements and Actions Rule apply? Yes, because S is about what the 2020 Business Tracker reported.

What are ALL elements of S\_restated?

[

"The 2020 Business Tracker reported that MiniCorp's success with its new line of eco-friendly products has inspired its competitors to invest in sustainable packaging -> contains verifiable information",

"The 2020 Business Tracker reported that MiniCorp's success with its new line of eco-friendly products has inspired its competitors to invest in renewable energy sources -> contains verifiable information",

"The 2020 Business Tracker reported that TurboCorp is an example of a competitor of MiniCorp that has been inspired by MiniCorp's success with its new line of eco-friendly products to invest in sustainable packaging -> contains verifiable information",

"The 2020 Business Tracker reported that TurboCorp is an example of a competitor of MiniCorp that has been inspired by MiniCorp's success with its new line of eco-friendly products to invest in renewable energy sources -> contains verifiable information",

"The 2020 Business Tracker reported that MegaCorp is an example of a competitor of MiniCorp that has been inspired by MiniCorp's success with its new line of eco-friendly products to invest in sustainable packaging -> contains verifiable information",

"The 2020 Business Tracker reported that MegaCorp is an example of a competitor of MiniCorp that has been inspired by MiniCorp's success with its new line of eco-friendly products to invest in renewable energy sources -> contains verifiable information",

"This highlights the ripple effect of sustainable business practices -> it's a generic statement, so it does not contain verifiable information",

]

### Example 2

Question: Who are key figures in the corporate sustainability movement?

Excerpt from response: There are also ongoing efforts to use partnerships as a means to improve sustainability, as demonstrated by Jane Smith. Many notable sustainability leaders like Smith do not work directly for a corporation, but her organization CleanTech has powerful partnerships with technology companies (e.g., MiniMax) to significantly improve waste management, demonstrating the power of collaboration.

Sentence of interest: Many notable sustainability leaders like Smith do not work directly for a corporation, but her organization CleanTech has powerful partnerships with technology companies (e.g., MiniMax) to significantly improve waste management, demonstrating the power of collaboration.

### Element Extraction System Prompt (Continued)

```
S = Many notable sustainability leaders like Smith do not work directly for a corporation, but her organization CleanTech has powerful partnerships with technology companies (e.g., MiniMax) to significantly improve waste management, demonstrating the power of collaboration.
Are there any clarifications needed to understand S based on its context?
"Smith" refers to Jane Smith.
S_restated = Jane Smith is an example of a notable sustainability leader who does not work directly for a corporation, but her organization CleanTech has powerful partnerships with technology companies, including MiniMax, to significantly improve waste management, which demonstrates the power of collaboration.
Does the Statements and Actions Rule apply? No.
What are ALL elements of S_restated?
[
"Jane Smith is an example of a notable sustainability leader -> 'notable' is not verifiable, but the rest is verifiable, so it contains verifiable information",
"Jane Smith does not work directly for a corporation -> contains verifiable information",
"Jane Smith has an organization called CleanTech -> contains verifiable information",
"CleanTech has powerful partnerships with technology companies to significantly improve waste management -> 'powerful' and 'significantly' are not verifiable, but the rest is verifiable, so it contains verifiable information",
"MiniMax is a technology company -> contains verifiable information",
"CleanTech has a partnership with MiniMax to significantly improve waste management -> 'significantly' is not verifiable, but the rest is verifiable, so it contains verifiable information",
"CleanTech demonstrates the power of collaboration -> it's an interpretation, so it does not contain verifiable information",
]
```

### Element Extraction User Prompt

```
Question:
{question}

Excerpt from response:
{excerpt}

Sentence of interest:
{sentence}
```

## N.2.3 Element Coverage

### Element Coverage System Prompt

```
## Overview
You will be given a question and an excerpt from the response to the question. You will also be given a dictionary of claims extracted from the excerpt (which will be referred to as C), and a dictionary of elements (which will be referred to as E).

An element is "covered by" a claim if the element is explicitly stated or strongly implied by the claim. For each element in E, your task is to determine whether the information in the element is covered by C by following these steps:
1. Print "E<insert number here>: <insert element here EXACTLY as written>" where number is the key in the dictionary and element is the value.
2. Determine whether the information in the element is covered by C. If the element has a note that some information is not verifiable, ignore that part and focus on the verifiable information. You CANNOT use any external information (e.g., if the element says "Politicians like John frequently discuss the economy" and C says "John frequently discusses the economy" but there is no claim that John is a politician, even if you have external information that John is a politician, the element is not fully covered by C). If C is more specific than E, you must check whether the specificity is merited based on the question
```

## Element Coverage System Prompt (Continued)

and the excerpt (i.e., if the elements should be more specific based on the context); if it is merited, then the element is fully covered by C. Print either "fully covered by C" or "not fully covered by C".

3. Repeat this process for all elements in E.

If the element is something like "John found X", "John reported X", "John emphasizes X", etc. (where John can be replaced with any entity or entities), it should be interpreted as a statement about what John says or does. For example, if the element is "John highlights that transparent communication is a critical part of Project Alpha", the claim "transparent communication is a critical part of Project Alpha" does not cover the element because it's missing the critical context that this is something John highlights. Let's call this the Statements and Actions Rule.

## Examples

### Example 1

Question: What are the key factors driving the shift towards sustainability in the fashion industry?

Excerpt from response: The growing public awareness of climate change has led to a surge in demand for sustainable fashion products. For example, MiniCorp recently launched a new line of eco-friendly scarves that have been well-received by consumers. The 2020 Business Tracker reported that this inspired its competitors, such as TurboCorp, to invest in sustainable packaging, highlighting the ripple effect of sustainable business practices.

Claims (C): {

- 1: "The 2020 Business Tracker reported that MiniCorp inspired its competitors to invest in sustainable packaging",
- 2: "The 2020 Business Tracker reported that TurboCorp was inspired by MiniCorp",
- 3: "TurboCorp is a competitor of MiniCorp",
- 4: "TurboCorp invested in sustainable packaging because it was inspired by MiniCorp",
- 5: "MiniCorp inspiring its competitors to adopt sustainable practice illustrates the ripple effect of sustainable business practices in the fashion industry",

}

Elements (E): {

- 1: "The 2020 Business Tracker reported that MiniCorp's success with its new line of eco-friendly scarves has inspired its competitors to invest in sustainable packaging",
- 2: "The 2020 Business Tracker reported that TurboCorp is an example of a competitor of MiniCorp that has been inspired by MiniCorp's success with its new line of eco-friendly scarves to invest in sustainable packaging",
- 3: "This highlights the ripple effect of sustainable business practices",

}

E1: The 2020 Business Tracker reported that MiniCorp's success with its new line of eco-friendly scarves has inspired its competitors to invest in sustainable packaging

- The Statements and Actions Rule applies because the element is about what the 2020 Business Tracker reported

- C1 says "The 2020 Business Tracker reported that MiniCorp inspired its competitors to invest in sustainable packaging"

- What is not explicitly stated or strongly implied by C, and is therefore grounds for lack of full coverage? It does not specify that it was MiniCorp's success with its new line of eco-friendly products that inspired its competitors. Therefore E1 is not fully covered by C

E2: The 2020 Business Tracker reported that TurboCorp is an example of a competitor of MiniCorp that has been inspired by MiniCorp's success with its new line of eco-friendly products to invest in sustainable packaging

- The Statements and Actions Rule applies because the element is about what the 2020 Business Tracker reported

- C2 says "The 2020 Business Tracker reported that TurboCorp was inspired by MiniCorp" and C3 says "TurboCorp is a competitor of MiniCorp" and C4 says "TurboCorp invested in sustainable packaging because it was inspired by MiniCorp"

- What is not explicitly stated or strongly implied by C?, and is therefore grounds for lack of full coverage Only C2 explicitly states that it was reported

## Element Coverage System Prompt (Continued)

by the 2020 Business Tracker, and C does not specify that it was MiniCorp's success with its new line of eco-friendly products that inspired TurboCorp. Therefore E2 is not fully covered by C

E3: This highlights the ripple effect of sustainable business practices  
- C5 says "MiniCorp inspiring its competitors to adopt sustainable practice illustrates the ripple effect of sustainable business practices in the fashion industry"

- What is not explicitly stated or strongly implied by C, and is therefore grounds for lack of full coverage? C5 is more specific than E3, so we need to check whether the specificity is merited based on the question and the excerpt. The question is about the key factors driving the shift towards sustainability in the fashion industry, and the excerpt discusses MiniCorp's success with its new line of eco-friendly scarves, so this specificity is merited. Therefore E3 is fully covered by C

### ### Example 2

Question: Who are key figures in the corporate sustainability movement?

Excerpt from response: There are also ongoing efforts to use partnerships as a means to improve sustainability, as demonstrated by Jane Smith. Many notable sustainability leaders like Smith do not work directly for a corporation, but her organization CleanTech has powerful partnerships with technology companies (e.g., MiniMax) to significantly improve waste management, demonstrating the power of collaboration.

Claims (C): {

- 1: "Jane is a sustainability leader",
- 2: "Jane doesn't work directly for a corporation",
- 3: "CleanTech has partnerships with technology companies to improve waste management",
- 4: "CleanTech has a partnership with MiniMax",

Elements (E): {

- 1: "Jane Smith is an example of a notable sustainability leader [note: 'notable' is not verifiable, but the rest is verifiable]",
- 2: "Jane Smith does not work directly for a corporation",
- 3: "Jane Smith has an organization called CleanTech",
- 4: "CleanTech has powerful partnerships with technology companies to significantly improve waste management [note: 'powerful' and 'significantly' are not verifiable, but the rest is verifiable]",
- 5: "MiniMax is a technology company",
- 6: "CleanTech has a partnership with MiniMax to significantly improve waste management [note: 'significantly' is not verifiable, but the rest is verifiable]",
- 7: "CleanTech demonstrates the power of collaboration",

Element 1: Jane Smith is an example of a notable sustainability leader

- C1 says "Jane is a sustainability leader", and "notable" is not verifiable so it can be ignored

- What is not explicitly stated or strongly implied by C, and is therefore grounds for lack of full coverage? Nothing. The verifiable parts of the element are explicitly stated. Therefore E1 is fully covered by C

Element 2: Jane Smith does not work directly for a corporation

- C2 says "Jane does not work directly for a corporation"

- What is not explicitly stated or strongly implied by C, and is therefore grounds for lack of full coverage? Nothing. The element is explicitly stated. Therefore E2 is fully covered by C

Element 3: Jane Smith has an organization called CleanTech

- C does not state that CleanTech is Jane's organization

- What is not explicitly stated or strongly implied by C, and is therefore grounds for lack of full coverage? Nothing. The element is explicitly stated. Therefore E3 is not fully covered by C

Element 4: CleanTech has powerful partnerships with technology companies to significantly improve waste management

- C3 says "CleanTech has partnerships with technology companies to improve waste management", and "powerful" and "significantly" are not verifiable so they can be ignored

### Element Coverage System Prompt (Continued)

- What is not explicitly stated or strongly implied by C, and is therefore grounds for lack of full coverage? Nothing. The verifiable parts of the element are explicitly stated. Therefore E4 is fully covered by C

Element 5: MiniMax is a technology company

- C4 says "CleanTech has a partnership with MiniMax"

- What is not explicitly stated or strongly implied by C, and is therefore grounds for lack of full coverage? It does not say that MiniMax is a technology company. Therefore E5 is not fully covered by C

Element 6: CleanTech has a partnership with MiniMax to significantly improve waste management

- C4 says "CleanTech has a partnership with MiniMax"

- What is not explicitly stated or strongly implied by C, and is therefore grounds for lack of full coverage? It does not say that the purpose of the partnership is to improve waste management. Therefore E6 is not fully covered by C

Element 7: CleanTech demonstrates the power of collaboration

- C3 says "CleanTech has partnerships with technology companies to improve waste management"

- What is not explicitly stated or strongly implied by C, and is therefore grounds for lack of full coverage? It does not explicitly state that CleanTech demonstrates the power of collaboration, but C strongly implies it. Therefore it is implied that E7 is fully covered by C

### Element Coverage User Prompt

Question (context for E):  
{question}

Excerpt from response (context for E):  
{excerpt}

Claims (C):  
{claims}

Elements (E):  
{elements}

## N.2.4 Decontextualization

### Decontextualization System Prompt

## Overview

You will be given a question, an excerpt from the response to the question, a sentence of interest from the excerpt, the claims derived from the sentence, and a claim of interest (which will be referred to as C).

A claim is "decontextualized" if (1) it is fully self-contained, meaning it can be understood in isolation (i.e., without the question, the excerpt, the sentence, and the other claims), AND (2) its meaning in isolation matches its meaning when interpreted alongside the question, the excerpt, the sentence, and the other claims.

Your task is to create C\_max, the maximally decontextualized version of C by following these steps:

1. Print "C = <insert claim of interest here EXACTLY as written>"
2. If someone read C without any context, would they have any questions? If yes, are any of these questions answered by the sentence or its context? If the reader would not have any questions, or none of the questions are clearly answered by the sentence or its context, or the claim already provides sufficient answers to the questions, print "C\_max = C". Otherwise, set C\_max equal to the maximally decontextualized claim that clarifies the answers to the questions that would be asked. Only include clarifications that are clearly attributable to the sentence and its context.

## Examples

## Decontextualization System Prompt (Continued)

### ### Example 1

Question: Provide an overview of the Supreme Court's importance in the United States.

Excerpt from response: Another example of the court's importance is the decision in Roe v. Wade in January 1973. It significantly affected abortion laws across the United States.

Sentence of interest: It significantly affected abortion laws across the United States.

All claims: ["The court's decision affected abortion laws across the United States."]

Claim of interest: The court's decision affected abortion laws across the United States.

C = The court's decision affected abortion laws across the United States.

Would someone reading C without any context have questions? They would likely ask which court made the decision and what the decision was.

Are any of these questions answered by the sentence or its context? The court is the Supreme Court, and the decision was Roe v. Wade in January 1973.

C\_max = The Supreme Court's decision in Roe v. Wade in January 1973 affected abortion laws across the United States.

### ### Example 2

Question: Who is Jane Doe?

Excerpt from response: Jane Doe has been a long-time advocate for environmental causes. In 2022, she spoke at Climate Action Now. She plans to speak at Youth for Climate, Moving Forward, and several other conferences next year.

Sentence of interest: She plans to speak at Youth for Climate, Moving Forward, and several other conferences next year.

All claims: ["She plans to speak at Youth for Climate next year", "She plans to speak at Moving Forward next year", "She plans to speak at other conferences next year"]

Claim of interest: She plans to speak at other conferences next year.

C = She plans to speak at other conferences next year.

Would someone reading C without any context have questions? They would likely ask who she is, what "other conferences" means, and what year "next year" refers to.

Are any of these questions answered by the sentence or its context? She is Jane Doe, "other conferences" means conferences other than Youth for Climate and Moving Forward (since they are covered by the other claims), and "next year" is 2023.

C\_max = Jane Doe plans to speak at conferences other than Youth for Climate and Moving Forward in 2023.

## Decontextualization User Prompt

Question:  
{question}

Excerpt from response:  
{excerpt}

Sentence of interest:  
{sentence}

All claims:  
{claims}

Claim of interest:  
{claim}



## N.3 Invalid Statements

### N.3.1 Invalid Sentences

#### Invalid Sentences System Prompt

##### ## Overview

You will be given a question, an excerpt from the response to the question, and a sentence of interest from the excerpt (which will be referred to as S).

Your task is to determine whether S, in light of its context, can be interpreted as a complete, declarative sentence by following these steps:

1. Print "S = <insert sentence of interest here EXACTLY as written>"
2. Describe the context for S.
3. Can S be interpreted as a complete, declarative sentence as is? If not, given its context, can it be rewritten as a complete, declarative sentence? If yes, print "S can be interpreted as a complete, declarative sentence"; otherwise, print "S cannot be interpreted as a complete, declarative sentence".

##### ## Examples

###### ### Example 1

Question: How can companies improve their sustainability practices?

Excerpt from response: Some examples include:

- Reducing energy consumption by using energy-efficient appliances
- Implementing recycling programs
- Sourcing materials from sustainable suppliers

Sentence of interest: Some examples include:

S = Some examples include:

Describe the context for S. The question is about how companies can improve their sustainability practices, and S is the header for a list of examples. Can S be interpreted as a complete, declarative sentence as is? If not, given its context, can it be rewritten as a complete, declarative sentence? S is not a complete, declarative sentence as is. Since it merely introduces the list of examples without providing a complete thought, it cannot be rewritten as a complete, declarative sentence. Therefore, S cannot be interpreted as a complete, declarative sentence.

###### ### Example 2

Question: How are companies improving their sustainability practices?

Excerpt from response: Some examples include:

- Reducing energy consumption by using energy-efficient appliances
- Implementing recycling programs
- Sourcing materials from sustainable suppliers

Sentence of interest: - Sourcing materials from sustainable suppliers

S = - Sourcing materials from sustainable suppliers

Describe the context for S. The question is about how companies are improving their sustainability practices, and the excerpt provides a list of examples. Can S be interpreted as a complete, declarative sentence as is? If not, given its context, can it be rewritten as a complete, declarative sentence? S is not a complete, declarative sentence as is. However, in the context of the question and the excerpt, it can be rewritten as "An example of how companies are improving their sustainability practices is sourcing materials from sustainable suppliers". Therefore, S can be interpreted as a complete, declarative sentence.

#### Invalid Sentences User Prompt

Question:  
{question}

Excerpt from response:  
{excerpt}

Sentence of interest:  
{sentence}

## N.3.2 Invalid Claims

### Invalid Claims System Prompt

```
## Overview
You will be given a claim (which will be referred to as C). Your task is to
determine whether C, in isolation, is a complete, declarative sentence, by
following these steps:
1. Print "C = <insert claim of interest here EXACTLY as written>"
2. In isolation, is C a complete, declarative sentence? After your reasoning,
print either "C is not a complete, declarative sentence" or "C is a complete,
declarative sentence".

## Examples
### Example 1
Claim: Sourcing materials from sustainable suppliers is an example of how
companies are improving their sustainability practices

C = Sourcing materials from sustainable suppliers is an example of how companies
are improving their sustainability practices
In isolation, is C a complete, declarative sentence? Yes, C is a complete,
declarative sentence.

### Example 2
Claim: Sourcing materials from sustainable suppliers

C = Sourcing materials from sustainable suppliers
In isolation, is C a complete, declarative sentence? It's missing a subject and
a verb, so C is not a complete, declarative sentence.
```

### Invalid Claims User Prompt

```
Claim:
{claim}
```