

FENICE: Factuality Evaluation of Summarization Based on Natural Language Inference and Claim Extraction

Alessandro Scire^{1,2}

Karim Ghonim²

Roberto Navigli²

Babelscape, Italy
¹scire@babelscape.com

Sapienza University of Rome
²{lastname}@diag.uniroma1.it

Abstract

Recent advancements in text summarization, particularly with the advent of Large Language Models (LLMs), have shown remarkable performance. However, a notable challenge persists as a substantial number of automatically-generated summaries exhibit factual inconsistencies, such as hallucinations. In response to this issue, various approaches for the evaluation of consistency for summarization have emerged. Yet, these newly-introduced metrics face several limitations, including lack of interpretability, focus on short document summaries (e.g., news articles), and computational impracticality, especially for LLM-based metrics. To address these shortcomings, we propose Factuality Evaluation of summarization based on Natural language Inference and Claim Extraction (FENICE), a more interpretable and efficient factuality-oriented metric. FENICE leverages an NLI-based alignment between information in the source document and a set of atomic facts, referred to as *claims*, extracted from the summary. Our metric sets a new state of the art on AGGREFACT, the de-facto benchmark for factuality evaluation. Moreover, we extend our evaluation to a more challenging setting by conducting a human annotation process of long-form summarization. In the hope of fostering research in summarization factuality evaluation, we release the code of our metric and our factuality annotations of long-form summarization at <https://github.com/Babelscape/FENICE>.

1 Introduction

In recent years, Natural Language Generation (NLG) approaches have achieved impressive results in many areas of Natural Language Processing (NLP), including Machine Translation (MT), and Question Answering (QA). When it comes to Text Summarization, the landscape differs somewhat. Despite the high performance boasted by conventional metrics such as ROUGE (Lin, 2004), several empirical investigations (Maynez et al., 2020;

Durmus et al., 2020; Kryscinski et al., 2020) revealed that a significant portion of automatically-generated summaries contains factual inconsistencies, e.g., hallucinations. Addressing this issue is crucial, as it stands as a barrier preventing the widespread application of such systems in real-world scenarios.

With the emergence of Large Language Models (LLMs) the challenge still persists, as their outputs continue to exhibit inconsistencies (Tam et al., 2022). What is concerning is that the heightened fluency boasted by these models (Wang et al., 2023b) may inadvertently lead inexperienced users to place excessive trust in their outputs, thereby fostering the spread of misinformation. For these reasons, the introduction of reliable automatic metrics that detect factual inaccuracies in summaries is becoming increasingly urgent.

As a response to this need, various factuality metrics have emerged, leveraging methodologies rooted in Question Answering (Scialom et al., 2021; Fabbri et al., 2022; Wang et al., 2020), Natural Language Inference (Zha et al., 2023; Chen and Eger, 2023; Laban et al., 2022, NLI), and semantic representations (Ribeiro et al., 2022; Goyal and Durrett, 2020; Fan et al., 2023). The latest trend in this domain involves works (Fu et al., 2023; Liu et al., 2023a; Gao et al., 2023) that introduce prompt engineering approaches using LLMs.

However, current factuality metrics suffer from one or more of the following drawbacks: i) they often assign a broad score to the entire summary without specifying which parts are factual or hallucinated, thereby lacking interpretability; ii) they are mainly designed for evaluating summaries of short documents, such as news articles, and iii) they are impractical from a computational standpoint, specifically in the case of LLM-based metrics.

To overcome these limitations, we introduce Factuality Evaluation of summarization based on Natural language Inference and Claim Extraction

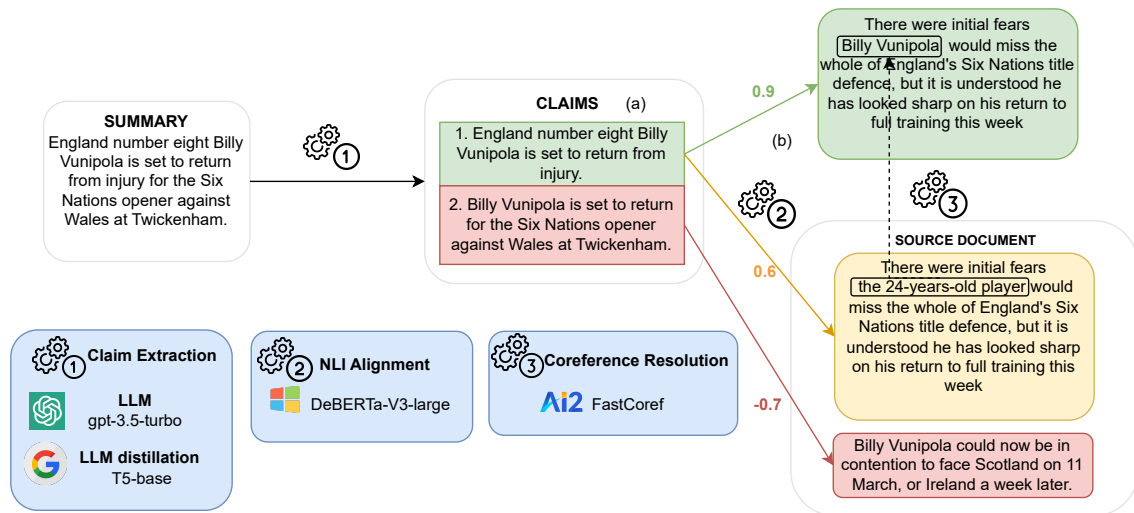


Figure 1: Overview of FENICE: the process begins with the extraction of claims from a given summary (step 1). Extracted claims are then aligned with specific sections of the input document (step 2). Finally, we refine the obtained alignments through a coreference-resolution-based approach (step 3). Best seen in color.

(FENICE), a state-of-the-art factuality metric that is more interpretable and efficient. As depicted in Figure 1, FENICE evaluates the factuality of a summary by leveraging an NLI-based alignment between information in the source document and a collection of atomic facts, referred to as *claims*, extracted from the summary. Our metric aligns claims from the summary with text from the input at three different levels of granularity: sentence, paragraph, and document. This novel alignment strategy facilitates interpretability: when presented with a claim, the metric user can either verify which specific part of the input document entails it, or, in the presence of hallucinations, they can observe that no alignment is provided by the metric.

Not only does our approach enhance interpretability but it also yields a positive impact on overall performance: FENICE achieves the highest average result in the AggreFact benchmark (Tang et al., 2023), the de-facto benchmark for summarization factuality evaluation. Furthermore, we evaluate the performance of our metric in a more challenging context, specifically targeting long-form summarization. To accomplish this, we conduct a human annotation process aimed at assessing long-form summarization factuality. When confronting the metrics outputs with the human labels, FENICE exhibits higher accuracy compared to its competitors, paving the way for future approaches for long document summarization evaluation.

In summary, our contributions are the following:

- We introduce FENICE, a state-of-the-art factuality metric for summarization evaluation;
- The alignments provided by our metric facilitate interpretability, highlighting the input document sections that are useful for summary claims verification;
- FENICE is designed to be computationally efficient, utilizing approximately 700 million parameters, thereby ensuring its applicability in real-world applications.
- We show the effectiveness of FENICE in the evaluation of long-form summarization factuality through a human annotation process.
- In the hope of fostering research in summarization factuality evaluation, we release the code of our metric and our human annotations at <https://github.com/Babelscape/FENICE>.

2 Related Work

NLI-based metrics. Drawing upon recent developments in the NLI field, several studies have introduced novel factuality metrics that exhibit stronger correlations with human judgments. Chen and Eger (2023) proposed a series of NLI metrics tailored for both MT and summarization evaluation. They conducted experiments in both reference-based and reference-less settings. In the latter setting, which serves as the standard for consistency evaluation,

they utilized the input document as the premise for the NLI system and the entire summary as the hypothesis for verification. While this approach correlates well with human consistency judgments, it lacks interpretability: on the one hand, when a summary is inconsistent, the end user is not provided with the parts of it that are deemed non-factual; on the other hand, when a summary is factually accurate, the metric provides no indication of the specific sections in the document that entail it.

In the same context, AlignScore (Zha et al., 2023) provides a novel alignment system, which is trained on several verification tasks, including NLI. While it achieves high performance in factuality evaluation, AlignScore also falls short in terms of interpretability, providing an overall score for the entire summary, which prevents users from identifying specific factual or non-factual spans within it. Moreover, it aligns the summary with large blocks of text from the input document, making it hard for users to gain insights from the output.

The Dependency Arc Entailment (Goyal and Durrett, 2020, DAE) metrics leverage the entailment between the dependency parse trees of document sentences and the summary. However, the highly abstractive nature of summaries, often encapsulating multiple facts in a single sentence, introduces a drawback, making their sentence-level comparison suboptimal.

Despite the significant progress achieved so far, we posit that there is still considerable room for enhancing NLI-based metrics, encompassing improvements in both performance and interpretability.

LLM-based metrics. Several studies (Shen et al., 2023; Chen et al., 2023; Wang et al., 2023a) have recently analyzed the efficacy of LLMs as summarization evaluators in multiple evaluation settings. Their experiments demonstrate that, while LLMs outperform existing automatic metrics, they are not yet ready to be considered reliable alternatives for human evaluation. Starting from these findings, TrueTeacher (Gekhman et al., 2023) analyzed the LLMs’ capability of generating large-scale factuality datasets and then trained smaller student models on such data, outperforming existing NLI metrics.

Another trend of approaches revolves around prompt engineering strategies (Luo et al., 2023; Fu et al., 2023; Liu et al., 2023a; Gao et al., 2023), where LLMs are typically instructed with the evaluation task description, the input document, and

the system-generated summary. Specifically, Fu et al. (2023) demonstrated the capabilities of such models in achieving multi-aspect, customized, and training-free evaluation. Relatedly, the recently-introduced G-Eval (Liu et al., 2023a) experimented with chain-of-thoughts (CoT) and form-filling instruction paradigms. They showed that G-Eval, configured with GPT-4, outperforms all previous methods in terms of correlation with the human consistency annotations in SummEval (Fabbri et al., 2021).

While achieving high correlations with human judgments, LLM-based metrics raise concerns about their real-world applicability. Firstly, many of these studies rely on closed-source LLMs, and the lack of information about the composition of their training dataset raises uncertainty: there is no guarantee that the test examples used in these studies are actually unseen at training time for these models. Additionally, their exceptionally-high computational cost poses a significant constraint on their practical deployment and hampers their reproducibility for research purposes.

Our approach differs from previous ones in that it utilizes an LLM solely for the task of claim extraction, rather than also utilizing it for generating the factuality annotations, which it may have already seen during training. Furthermore, we provide a distilled version of this LLM—a more compact generative Transformer—trained to emulate the output of the LLM, which ensures the practicality and reproducibility of our approach.

Claim extraction-based evaluation. A growing number of works explores the employment of claim extraction for the purpose of evaluating generated text. Among these, Min et al. (2023) proposed a method for evaluating long-form text generation, through the validation of their extracted claims given a reliable knowledge source, e.g., Wikipedia.

In the context of summarization evaluation, Nenkova et al. (2007) introduced Pyramid, the first manual annotation protocol in this context. However, the scope of Pyramid and its automatic approximations (Gao et al., 2019; Harnly et al., 2005) is to evaluate summarization relevance, which differs significantly from the factuality-oriented focus of our study.

To the best of our knowledge, we are the first to consider claims as the foundation for a factuality-oriented summarization metric.

3 FENICE

We now provide an overview of the FENICE algorithm. Initially, we extract a list of atomic facts, referred to as claims, from a given summary (Section 3.1). Subsequently, we detail the process of obtaining a factuality score for each extracted claim. This involves leveraging NLI, with the claim serving as the hypothesis, and the document sentences (Section 3.2) or longer blocks of text at different levels of granularity (Section 3.4) serving as premises. In Section 3.3, we elaborate on a refinement of the obtained scores through coreference resolution.

3.1 Claim extraction

Following the Atomic Content Unit (ACU) definition outlined in Liu et al. (2023b), we use the term *claim* to denote an *atomic fact*, i.e., an elementary information unit found in a summary, that does not require further subdivision. Figure 1(a) depicts an example list of claims extracted from a summary.

We frame the claim extraction problem as an end-to-end autoregressive generation task. Given a prompt consisting of the task description and the summary to be processed, the model generates a sequence of claims expressed within the provided summary. To guide the model in meeting the desired output format and enhancing its task comprehension, we include an example of an input summary and the expected claims in the prompt (see Appendix A).

We opt to center our evaluation strategy around claims rather than sentences, primarily because of their enhanced interpretability and their atomic nature, which facilitates alignment with information from the input.

3.2 NLI-based claim scoring

The NLI formulation lies at the core of our novel claim scoring algorithm. NLI is a task that involves determining the logical relationship between two text segments: a *premise* and a *hypothesis*. Formally, given a premise *pre* and a hypothesis *hyp*:

$$\text{NLI}(pre, hyp) = (p_{ent}, p_{neu}, p_{con}) \quad (1)$$

where p_{ent} , p_{neu} , and p_{con} denote the respective probabilities of *pre* entailing, being neutral about, or contradicting *hyp*.

Starting from Equation 1, we define NLIScore as the difference between the entailment and contradiction probabilities, obtained using the document

sentence as premise and the claim extracted from the summary as hypothesis:

$$\text{NLIScore}(s, c) = p_{ent} - p_{con}$$

where the premise s is a sentence of the input document, and the hypothesis c is a claim extracted from the summary. NLIScore measures the degree to which a specific premise can be utilized to verify the claim c . The inclusion of the contradiction probability serves the purpose of penalizing claims that are in conflict with any content in the source document.

Finally, we introduce FENICE_{sent} as the function that assigns a score to a given claim c by considering the sentences in the source document D as premises:

$$\text{FENICE}_{sent}(D, c) = \max_{s \in D} \text{NLIScore}(s, c) \quad (2)$$

We also define align_{sent} as the corresponding alignment function:

$$\text{align}_{sent}(D, c) = \operatorname{argmax}_{s \in D} \text{NLIScore}(s, c) \quad (3)$$

For each claim c from the summary, this establishes a one-to-one alignment between c and the highest-scoring sentence from the input document.

3.3 Coreference Resolution

Importantly, given a claim c and its aligned sentence $s = \text{align}_{sent}(D, c)$, $\text{NLIScore}(s, c)$ may not reflect the true degree of entailment for the (s, c) pair due to an entity or concept being mentioned in the two texts with distinct surface forms. For instance, consider the claim in Figure 1(a): "England number eight Billy Vunipola is set to return from injury" and its aligned source document sentence: "There were initial fears the 24-years-old player would miss [...]". In this case, the NLI model assigns only a moderate score of 0.6 to this pair due to the limited context, making it challenging to establish a clear correspondence between "Billy Vunipola" and "the 24-years-old player".

To tackle this challenge, we apply a Coreference Resolution model on the source document, which identifies clusters of mentions that refer to the same entity or concept. We then introduce a coreference-resolution-aware version of Equation 2 as follows:

$$\begin{aligned} s^* &= \text{align}_{sent}(D, c) \\ \text{FENICE}_{coref}(D, c) &= \max_{s' \in \text{Co}(s^*)} \text{NLIScore}(s', c) \end{aligned} \quad (4)$$

where $Co(s^*)$ is the set of sentences obtained by replacing each entity occurring in sentence s^* with all its coreferential mentions.¹

For instance, given "the 24-years-old player" we identify its coreferential spans in the document, such as "Billy Vunipola". Subsequently, we generate an alternative version of the original sentence by replacing "the 24-years-old player" with "Billy Vunipola", resulting in the sentence: "*There were initial fears Billy Vunipola would miss [...]*". As depicted in Figure 1(b), this approach facilitates the alignment process of the claim "*England number eight Billy Vunipola is set to return from injury*" to the revisited version of the sentence, which contains the same span for the given named entity.

3.4 Aligning claims across multiple input text granularities

Given the abstract nature of summaries, the verification of a claim may require a context that extends beyond a single sentence from the input. For example, to verify the claim: "*Regular exercise contributes to overall well-being*", the following input passage would be necessary: "*Studies have consistently shown that individuals who engage in regular physical activity experience improved mental health. Furthermore, regular exercise is linked to a reduced risk of chronic diseases and enhanced cardiovascular health.*".

Therefore, we define the following variant of Equation 2 for longer passages of k sentences:

$$FENICE_p(D, c, k) = \max_{i=1, \dots, |D|-k+1} \text{NLIScore}(s_i^{i+k-1}, c) \quad (5)$$

where s_i^{i+k-1} is the concatenation of s_i, \dots, s_{i+k-1} , i.e., the k consecutive sentences starting from sentence i in document D .

Our approach involves choosing k such that $k < |D|$, alongside considering the entire document as premise, i.e., $k = |D|$. Thus, our function for scoring claims given passages of multiple lengths from the source document is given by the following:

$$FENICE_{mul}(D, c, j) = \max\{FENICE_p(D, c, j), FENICE_p(D, c, |D|)\} \quad (6)$$

where j is the desired length for the contiguous blocks of the source sentences. Then, we define

¹We limit this augmentation to the sentence s^* , which is the result of Equation 3.

$FENICE_{cl}$ as the factuality score assigned to a claim c :

$$FENICE_{cl}(D, c, j, T) = \begin{cases} FENICE_{mul}(D, c, j) & \text{if } FENICE_{coref}(D, c) < T; \\ FENICE_{coref}(D, c) & \text{otherwise.} \end{cases} \quad (7)$$

where T is a threshold. We introduce T to limit the application of $FENICE_{mul}$ to only a subset of claims that requires further verification, i.e., the low-scoring ones. We set $j = 5$ and $T = 0.8$ in all our experiments (Section 4.2) as a result of tuning them on the validation splits.

Finally, given the set of extracted claims C from a summary S , and a source document D , we introduce FENICE as:

$$FENICE(D, C, j, T) = \frac{\sum_{c \in C} FENICE_{cl}(D, c, j, T)}{|C|} \quad (8)$$

which is the factual consistency score attributed by our metric to summary S .

4 Experiments and Results

4.1 Claim extraction

Experimental setup. As detailed in Section 3.1, we formulate the claim extraction problem as a sequence generation task. To accomplish this, we designed a prompt that includes the task description, an input-output example, and the summary to be processed. The prompt utilized is available in Appendix A. We experiment with *GPT-3.5-turbo* as claim extractor.² We additionally provide a distilled version of the claim extractor, which consists of a T5_{base} (Raffel et al., 2019) model trained on the LLM outputs. To create the training and evaluation datasets for the distilled model, we apply the LLM-based claim extractor on the reference summaries from two widely-used summarization datasets, namely XSum (Narayan et al., 2018) and CNN/DailyMail (Nallapati et al., 2016). We split the dataset obtained into training, validation, and test sets following a 80%-10%-10% proportion, comprising 430,946 summary-claims pairs for training, and 53,868 for both the validation and test partitions. We train our system on a single GPU *NVIDIA GeForce RTX 3090* for a total of 1M steps, using AdaFactor (Shazeer and Stern, 2018) as optimizer with learning rate of $1 \cdot 10^{-5}$.

²We use the *gpt-3.5-turbo* API, version dated 2nd December 2023.

Model	easiness _P	easiness _R	easiness _{F1}
GPT-3.5	80.1	70.9	74.9
T5 _{dist_GPT}	79.2	68.8	73.4

Table 1: Easiness Precision (easiness_P), Recall (easiness_R), and F1 score (easiness_{F1}) results for the LLM-based claim extractor, namely GPT-3.5, and its distilled version T5_{dist_GPT}, assessed on ROSE (Liu et al., 2023b).

Datasets. We evaluate our distilled version of the claim extractor on the test split of the LLM-generated claim extraction dataset. We additionally assess the performance of both the LLM-based claim extractor and our distilled version on ROSE (Liu et al., 2023b), a dataset that includes summaries and human-written claims among its annotations. While this dataset was designed to serve as the basis for a human annotation protocol for summarization evaluation, we utilize the released manually-extracted claims for assessing our systems’ performance.

Metrics. Given the extractive nature of this task, we consider ROUGE (Lin, 2004) as the basis for a suitable evaluation metric for claim extraction. Following (Zhang and Bansal, 2021) we define easiness_P as the score obtained by computing the ROUGE-1 metric for each pair of system-generated and gold standard claims, subsequently calculating the maximum, and averaging over the total number of generated claims:

$$\text{easiness}_P(S, H) = \frac{\sum_{c \in S} \max_{c^* \in H} \text{R1}(c, c^*)}{|S|} \quad (9)$$

where R1 is the ROUGE-1-F1 metric, and S, H are the sets of system- and human-written claims, respectively. We then introduce easiness_R as the recall-oriented version of Equation 9, obtained by averaging over the gold standard claims:

$$\text{easiness}_R(S, H) = \frac{\sum_{c^* \in H} \max_{c \in S} \text{R1}(c, c^*)}{|H|} \quad (10)$$

Results and discussion. In Table 1 we report the results of the easiness_P, and easiness_R metrics, along with the F1 version of Equations 9 and 10, i.e., easiness_{F1}, obtained on ROSE. In Appendix B we additionally report the results obtained when using BERTScore (Zhang et al., 2020) as backbone metric for our evaluation instead of ROUGE.

Model	easiness _P	easiness _R	easiness _{F1}
T5 _{dist_GPT}	89.7	89.9	89.5

Table 2: Easiness metric results of the distilled version of the LLM-based claim extractor, i.e., T5_{dist_GPT} evaluated on the test set of the dataset composed of claims generated by GPT-3.5.

As we can see from Table 1, our distilled version performs comparably with the original one based on GPT-3.5. These results indicate the effectiveness of the distillation process, ensuring that our model can be used reliably for extracting claims. To further evaluate the quality of the claims generated by our distilled model, we compare them to those generated by the original GPT-3.5 model, using the latter as references. The results of this evaluation are presented in Table 2, demonstrating the strong ability of our T5 model to replicate the claims generated by GPT-3.5.

4.2 Factuality Evaluation

Experimental setting. We leverage an off-the-shelf model trained on NLI to align the summary-extracted claims with text from the input document. Specifically, we employ a version of DeBERTa-v3-large (He et al., 2020) fine-tuned on multiple NLI datasets,³ comprising MultiNLI (Williams et al., 2017), LingNLI (Parrish et al., 2021), WANLI (Liu et al., 2022), and Fever-NLI, as well as Adversarial-NLI, both introduced in Nie et al. (2019). We adopt spaCy’s *en-core-web-sm* to segment the input documents into sentences, and F-Coref (Otmazgin et al., 2022) for the extraction of the coreferential spans (Section 3.3).

Evaluation datasets. We assess the effectiveness of FENICE on AGGREGFACT (Tang et al., 2023), the de-facto benchmark for evaluating summarization factuality metrics. AGGREGFACT comprises multiple datasets of news articles and manually annotated summaries with factuality labels. The benchmark is categorized into three splits based on the summarization systems employed: i) FTSOTA, which includes the outputs of state-of-the-art pre-trained summarizers, ii) EXFORMER, consisting of summaries of early Transformer-based summarization models, and iii) OLD, which comprises the

³We loaded the pre-trained model weights from the Hugging Face Hub, available at: <https://huggingface.co/MoritzLaurer/DeBERTa-v3-large-mnli-fever-anli-ling-wanli>

	Agg-CNN			Agg-XSum			AVG
	FTS	EXF	OLD	FTS	EXF	OLD	
Random Baseline	50.0	50.0	50.0	50.0	50.0	50.0	50.0
DAE*	59.4	67.9	69.7	73.1	-	-	67.5
QuestEval	63.7	64.3	65.2	61.6	60.1	59.7	63.5
SummaC-ZS	63.3	76.5	76.3	56.1	51.4	53.3	65.3
SummaC-Cv	70.3	69.8	78.9	67.0	64.6	67.5	70.5
TrueTeacher-11B	65.7	57.7	<u>81.9</u>	<u>75.2</u>	68.4	52.8	67.0
QAFactEval	61.6	69.1	80.3	65.9	59.6	60.5	67.0
MENLI	51.7	52.8	68.4	58.3	59.7	73.9	60.8
AlignScore	53.5	73.9	78.0	80.2	79.9	63.7	<u>71.5</u>
ChatGPT-ZS	66.2	64.5	74.3	62.6	69.2	60.1	65.7
ChatGPT-CoT	49.7	60.4	66.7	56.0	60.9	50.1	57.9
ChatGPT-DA	48.0	63.6	71.0	53.6	65.6	61.5	60.1
ChatGPT-Star	55.8	65.8	71.2	57.7	70.6	53.8	63.8
FENICE _{GPT_claims}	<u>68.2</u>	68.8	82.1	73.9	<u>73.5</u>	<u>69.9</u>	72.7
FENICE _{T5_claims}	61.6	67.8	80.6	71.4	70.7	67.7	70.0

Table 3: Balanced accuracy results on the test sets of the AGGREGFACT datasets. Following Tang et al. (2023) we exclude the performance of DAE on the EXFORMER (EXF in table) and OLD datasets in the AGGREGFACT-XSUM partition, since it was trained on XSumFaith (Goyal and Durrett, 2021) which is part of those splits. Results in **bold** indicate the best performance, while underlined values indicate the second best.

outputs of older approaches.

Baselines. We evaluate the performance of FENICE against several baselines that align with recent trends in the field, specifically focusing on NLI, QA, and LLM-based metrics.

NLI-based metrics. We evaluate several NLI-based metrics, such as AlignScore (Zha et al., 2023), MENLI (Chen and Eger, 2023), SummaC-ZS and SummaC-Conv, both introduced in (Laban et al., 2022), and DAE (Goyal and Durrett, 2020).

QA-based metrics. Question Answering-based factuality metrics typically depend on a Question Generation (QG) model designed to formulate questions based on the input document as context, and check whether information in the summary can be used to answer such questions, and vice versa. For our study, we choose the top-performing QA metric on AGGREGFACT as baseline, namely QAFactEval (Fabbri et al., 2022).

LLM-based metrics. We benchmark recent ChatGPT-based evaluation metrics: ChatGPT-ZS and ChatGPT-CoT (Luo et al., 2023) which leverage GPT-3.5 to generate a binary factuality label, in zero-shot and Chain-of-Thought paradigms, respectively. We also include ChatGPT-DA and ChatGPT-

Metric	CNN(FTS)	XSUM(FTS)	AVG
DAE	65.4±4.4	70.2±2.3	67.8
QuestEval	<u>70.2</u> ±3.2	59.5±2.7	64.9
SummaC-ZS	64.0±3.8	56.4±1.2	60.2
SummaC-Conv	61.0±3.9	65.0±2.2	63.0
QAFactEval	67.8 ±4.1	63.9±2.4	65.9
TrueTeacher-11B	62.0±1.3	74.9 ±1.2	68.4
MENLI	63.4±2.8	59.0±1.8	61.2
AlignScore	62.7±3.1	69.4±1.9	66.1
ChatGPT-ZS	56.3±2.9	62.7±1.7	59.5
ChatGPT-COT	52.5±3.3	55.9±2.1	54.2
ChatGPT-DA	53.7±3.5	54.9±1.9	54.3
ChatGPT-Star	56.3±3.1	57.8±0.2	57.1
FENICE _{GPT_claims}	70.5 ±1.8	<u>72.8</u> ±0.3	71.6
FENICE _{T5_claims}	67.7±3.0	70.1±1.8	<u>68.9</u>

Table 4: Balanced accuracy results on the AGGREGFACT-FTSOTA split obtained with the single-threshold setting. Results in **bold** indicate the best performance, while underlined values represent the second best.

Star (Wang et al., 2023a), which prompt LLMs to generate factuality scores on a scale of 0-100 and 1-5, respectively. We additionally include the results of the 11B-parameter T5 version of the TrueTeacher metric.

Results and discussion. We assess the performance of two distinct variants of our metric: FENICE_{GPT_claims}, which employs LLM-based claims, and FENICE_{T5_claims}, featuring the knowledge-distilled claim extractor. To convert FENICE scores into binary labels, we tune a different threshold on each validation split of the AGGREGFACT datasets (i.e., FTSOTA, EXFORMER, and OLD) and then compute the balanced accuracy on the corresponding test set, as indicated in Tang et al. (2023).

In Table 3, we present the balanced accuracy (Brodersen et al., 2010) scores on the AGGREGFACT test sets. FENICE_{GPT_claims} sets a new state-of-the-art performance on AGGREGFACT, resulting in the highest average balanced accuracy score.

Moreover, we specifically analyze the metric’s ability to evaluate the outputs of state-of-the-art models. Following the recommendation of Tang et al. (2023), we report the results of our metrics on AGGREGFACT-FTSOTA split in the single-threshold setting. As outlined by Tang et al. (2023), we fine-tune a threshold on the validation sets of the CNN-FTSOTA and XSUM-FTSOTA splits and subsequently apply it to convert our metric scores into

binary labels on the test set. From the outcomes presented in Table 4, we note that FENICE_{GPT_claims} obtains the best average results in this split, also due to its high performance on XSUM-FtSOTA. Notably, we consider this latter setting to be the most challenging for a metric, given that it comprises summaries of state-of-the-art systems trained on XSum references, which exhibit the highest level of abstraction. In this context, FENICE_{T5_claims} slightly outperforms TrueTeacher, despite utilizing significantly fewer parameters (700M vs 11B).

4.3 Ablation study

We conduct an analysis of the impact of our primary design choices, namely, claim extraction, coreference resolution, the alignment with multiple text granularities, and our backbone NLI model.

First, we introduce NLI_{sent}, a simple baseline leveraging NLI-based sentence-level alignment between the summary and the input document. Specifically, the factuality score for a summary sentence s' is obtained through the following variant of Equation 2, considering summary sentences as hypotheses instead of claims:

$$\text{NLI}_{sent}(D, s') = \max_{s \in D} \text{NLIScore}(s, s') \quad (11)$$

As outlined in Section 3.4, we obtain the factuality score for a summary by averaging the result of Equation 11 applied to its sentences.

We then assess the impact of focusing the factuality evaluation on claims extracted from the summary rather than its sentences. To achieve this, we introduce NLI_{claim}, a modification of the previous baseline that assigns factuality scores to the claims using Equation 2. Subsequently, we evaluate the performance of NLI_{coref}, derived by integrating the coreference resolution strategy (Section 3.3) into the previous baseline. Additionally, we assess the impact of including the alignment strategy with input text at multiple granularities (Section 3.4) by presenting the performance of FENICE. Table 5 displays the balanced accuracy scores of FENICE on the AGGREGATE-FtSOTA split, comparing them to the aforementioned baselines. The values in the ΔAVG column represent the differences in terms of average balanced accuracy resulting from the incremental inclusion of each component. The results indicate that each component contributes significantly to the final performance. The alignment strategy with input text at multiple granularity has the most substantial impact, resulting in a signifi-

Metric	CNN	XSUM	AVG	ΔAVG
NLI _{sent}	52.2±1.4	59.6±0.6	55.9	—
NLI _{claim}	54.3±2.1	63.2±0.8	58.8	+2.9
NLI _{coref}	59.7±2.3	63.7±0.9	61.5	+2.7
FENICE	70.5±3.0	72.8±0.3	71.6	+10.1

Table 5: Ablation study evaluating the impact of component additions on FENICE on the AGGREGATE-FtSOTA split. We report the results of the following baselines: i) NLI_{sent}, utilizing summary sentences instead of claims, ii) NLI_{claim}, which leverages claims, and iii) NLI_{coref}, obtained by integrating the coreference-resolution-based alignments.

cant gain of 10.1 points when included. Additionally, we evaluate the performance of FENICE utilizing RoBERTa as the NLI model for our alignments (FENICE_{RoBERTa}), instead of DeBERTa (Appendix C).

4.4 Long-form summarization evaluation

In recent years, there has been a noticeable shift in research focus away from summarizing short texts, predominantly news articles, to handling longer documents like scientific papers (Gupta et al., 2021) and books (Kryściński et al., 2022; Scirè et al., 2023). Despite this rapid evolution in the development of approaches and resources, the evaluation of summarization has been slower to adapt, with only a limited number of works, which are primarily aimed at defining the basic guidelines for this task (Krishna et al., 2023).

Human evaluation. As an initial step towards addressing this gap, we conducted an evaluation of our metric in the context of long-form text summarization. To do so, we curated a set of annotations for assessing summaries of lengthy texts.

Drawing upon Echoes from Alexandria (Scirè et al., 2023), a multilingual resource for book summarization, we selected 26 English narratives, each paired with a reference abstractive summary. The reference summaries were refined by removing irrelevant details such as the story title, the genre, and the author’s name. As an example, the initial summary *"Isaac Asimov’s 'Let’s Get Together’ is a science fiction novel which depicts the continuation of Cold War hostilities..."* underwent manual editing and was refined to *"The story depicts the continuation of Cold War hostilities..."*. To augment our selection, we incorporated summaries generated by 12 book summarization systems, as presented by Scirè et al. (2023), applied to our se-

Dataset	Avg. Sentences	Avg. Tokens
Ours	664.2	11,817.9
SummEval	18.0	363.1
SummaC	23.0	485.5

Table 6: Our human annotations for long-form text summarization statistics compared to standard-form (news) factuality benchmarks, i.e., SummaC and SummEval.

lected stories. The systems chosen were versions of BART (Lewis et al., 2020), LongT5 (Guo et al., 2022), and LED (Beltagy et al., 2020), fine-tuned on the *Echo-XSum* split of the *Echoes from Alexandria* dataset. This process resulted in a comprehensive set of 338 text-summary pairs.

Next, we tasked an expert annotator with reading all the stories and assigning a binary factuality label to each of the associated summaries.⁴ As the foundation for the annotation process, we considered a summary factual if its content was explicitly stated or implied by the source document; otherwise, it was labeled as non factual. Through this process, we acquired factuality labels for all the 338 text-summary pairs.

Moreover, in order to measure the inter-annotator agreement we asked a second expert to manually annotate a subset of 52 summaries. We then computed the Cohen’s kappa coefficient, obtaining a moderate agreement of 0.74. We deemed this outcome reasonable, particularly when considering the high complexity of assessing long-form summarization compared to the standard setting. Table 6 compares the length of source texts in our proposed long-form summarization evaluation dataset with those in existing datasets, specifically SummaC (Laban et al., 2022) and SummEval (Fabri et al., 2021). In our dataset, annotators were tasked with reading very long source texts, averaging 664 sentences, in order to assess the factuality of the provided summaries. The distribution of factual and unfactual labels in our dataset is detailed in Appendix D. To align with the evaluation setup of AGGREGFACT, we partitioned our annotations into test and validation sets using an 80/20 split.

Results and discussion. We follow the experimental setup outlined in Section 4.2. We employ

⁴We supplied the annotator with the PDF files containing the text of the stories. We paid the annotator according to the standard salary for their geographical location. The annotator took an average of 2 hours to read a story, and an average of 5 minutes to annotate the corresponding summaries.

Metric	Balanced Accuracy
AlignScore	61.3
MENLI	61.7
DAE	51.4
FENICE _{GPT_claims}	66.2
FENICE _{T5_claims}	<u>65.7</u>

Table 7: Balanced accuracy results of FENICE metrics and the selected baselines on our set of human annotations for long form summarization factuality. Results in **bold** indicate the best performance, while underlined values represent the second best.

AlignScore, MENLI, and DAE as baselines, considering their relatedness to our work. In Table 7, we present the balanced accuracy results on the test split of our dataset. The results obtained highlight that FENICE metrics outperform the baselines, with FENICE_{GPT_claims} obtaining the highest performance and FENICE_{T5_claims} coming second. Upon manual inspection of our alignments, we observe that a significant proportion of claims are aligned with source text having the length of a paragraph. We posit that this behaviour is due to the high abstractiveness of the stories’ summaries, which typically require contexts that are longer than a sentence in order to be verified. We hypothesize that our metric is well-suited for both standard and long-form summarization thanks to its adaptive capacity to choose the appropriate input text granularity for claim verification.

5 Conclusion and Future Work

We introduced FENICE, a novel metric designed for the factuality evaluation of summarization. By leveraging NLI-based alignments between claims extracted from the summary and the source text at multiple levels of granularities, FENICE facilitates interpretability. Our experiments demonstrate that our metric attains state-of-the-art performance in standard summarization evaluation datasets. Additionally, we showcase the effectiveness of our metric in a more-challenging setting, i.e., the evaluation of long-form summarization. To do so, we conducted a human annotation process for this task. Our work paves the way for future research focused on assessing long-form summarization factuality. As future work, we plan to improve our zero-shot alignment system’s performance with a tailored fine-tuning on summarization datasets.



Limitations

While our approach incorporates multiple components, including NLI, coreference resolution and claim extraction, in order to enhance the robustness and comprehensiveness of the factuality evaluation metric, it is essential to acknowledge that the involvement of these diverse systems introduces a layer of complexity and potential sources of error. The inherent variability and imperfections in each of these components may impact the overall performance and necessitate careful consideration in the interpretation of the results.

Our experiments are also constrained by the fact that the models employed are designed for working in English only. In order to extend this approach to other languages, the multilingual variants of the NLI, claim extractor, and coreference resolution models will need to be employed. However, the performance of such multilingual models needs a proper evaluation that we defer to future work.

The claim-level granularity of the outputs of FENICE, and the alignment of claims with portions of the input document it provides, ensures a higher level of interpretability compared to previous metrics. Yet, defining a quantitative evaluation for this aspect remains an ongoing challenge.

Acknowledgements

We gratefully acknowledge the support of the PNRR MUR project PE0000013-FAIR.  

We gratefully acknowledge the CREATIVE project (CRoss-modal understanding and gENERATION of Visual and tEXtual content), which is funded by the MUR Progetti di Ricerca di Rilevante Interesse Nazionale programme (PRIN 2020). Alessandro Scirè and Karim Ghonim conducted this work during their enrollment in the Italian National Doctorate on Artificial Intelligence at Sapienza University of Rome.

References

- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *CoRR*, abs/2004.05150.
- Kay Henning Brodersen, Cheng Soon Ong, Klaas Enno Stephan, and Joachim M. Buhmann. 2010. [The balanced accuracy and its posterior distribution](#). In *2010 20th International Conference on Pattern Recognition*, pages 3121–3124.
- Shiqi Chen, Siyang Gao, and Junxian He. 2023. [Evaluating factual consistency of summaries with large language models](#).
- Yanran Chen and Steffen Eger. 2023. [Menli: Robust evaluation metrics from natural language inference](#).
- Esin Durmus, He He, and Mona Diab. 2020. [FEQA: A question answering evaluation framework for faithfulness assessment in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5055–5070, Online. Association for Computational Linguistics.
- Alexander Fabbri, Chien-Sheng Wu, Wenhao Liu, and Caiming Xiong. 2022. [QAFactEval: Improved QA-based factual consistency evaluation for summarization](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2587–2601, Seattle, United States. Association for Computational Linguistics.
- Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. [SummEval: Re-evaluating summarization evaluation](#). *Transactions of the Association for Computational Linguistics*, 9:391–409.
- Jing Fan, Dennis Aumiller, and Michael Gertz. 2023. [Evaluating factual consistency of texts with semantic role labeling](#).
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. [Gptscore: Evaluate as you desire](#).
- Mingqi Gao, Jie Ruan, Renliang Sun, Xunjian Yin, Shiping Yang, and Xiaojun Wan. 2023. [Human-like summarization evaluation with chatgpt](#).
- YanJun Gao, Chen Sun, and Rebecca J. Passonneau. 2019. [Automated pyramid summarization evaluation](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 404–418, Hong Kong, China. Association for Computational Linguistics.
- Zorik Gekhman, Jonathan Herzig, Roei Aharoni, Chen Elkind, and Idan Szpektor. 2023. [Trueteacher: Learning factual consistency evaluation with large language models](#).
- 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.
- Tanya Goyal and Greg Durrett. 2021. [Annotating and modeling fine-grained factuality in summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1449–1462, Online. Association for Computational Linguistics.

- Mandy Guo, Joshua Ainslie, David Uthus, Santiago Ontanon, Jianmo Ni, Yun-Hsuan Sung, and Yinfei Yang. 2022. [LongT5: Efficient text-to-text transformer for long sequences](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 724–736, Seattle, United States. Association for Computational Linguistics.
- Vivek Gupta, Prerna Bharti, Pegah Nokhiz, and Harish Karnick. 2021. [SumPubMed: Summarization dataset of PubMed scientific articles](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: Student Research Workshop*, pages 292–303, Online. Association for Computational Linguistics.
- Aaron Harnly, Ani Nenkova, Rebecca Passonneau, and Owen Rambow. 2005. Automation of summary evaluation by the pyramid method. In *International Conference on Recent Advances in Natural Language Processing, RANLP 2005 - Proceedings*, International Conference Recent Advances in Natural Language Processing, RANLP, pages 226–232. Association for Computational Linguistics (ACL). International Conference on Recent Advances in Natural Language Processing, RANLP 2005 ; Conference date: 21-09-2005 Through 23-09-2005.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. [DeBERTa: Decoding-enhanced BERT with disentangled attention](#). *CoRR*, abs/2006.03654.
- Kalpesh Krishna, Erin Bransom, Bailey Kuehl, Mohit Iyyer, Pradeep Dasigi, Arman Cohan, and Kyle Lo. 2023. [LongEval: Guidelines for human evaluation of faithfulness in long-form summarization](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1650–1669, Dubrovnik, Croatia. Association for Computational Linguistics.
- Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. 2020. [Evaluating the factual consistency of abstractive text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9332–9346, Online. Association for Computational Linguistics.
- Wojciech Kryściński, Nazneen Rajani, Divyansh Agarwal, Caiming Xiong, and Dragomir Radev. 2022. [Booksum: A collection of datasets for long-form narrative summarization](#).
- Philippe Laban, Tobias Schnabel, Paul N. Bennett, and Marti A. Hearst. 2022. [SummaC: Re-visiting NLI-based models for inconsistency detection in summarization](#). *Transactions of the Association for Computational Linguistics*, 10:163–177.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Alisa Liu, Swabha Swayamdipta, Noah A. Smith, and Yejin Choi. 2022. [WANLI: Worker and AI collaboration for natural language inference dataset creation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 6826–6847, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023a. [G-eval: Nlg evaluation using gpt-4 with better human alignment](#).
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Yixin Liu, Alexander R. Fabbri, Pengfei Liu, Yilun Zhao, Linyong Nan, Ruilin Han, Simeng Han, Shafiq Joty, Chien-Sheng Wu, Caiming Xiong, and Dragomir Radev. 2023b. [Revisiting the gold standard: Grounding summarization evaluation with robust human evaluation](#).
- Zheheng Luo, Qianqian Xie, and Sophia Ananiadou. 2023. [Chatgpt as a factual inconsistency evaluator for text summarization](#).
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. [On faithfulness and factuality in abstractive summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1906–1919, Online. Association for Computational Linguistics.
- Sewon Min, Kalpesh Krishna, Xinxin Lyu, Mike Lewis, Wen tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [Factscore: Fine-grained atomic evaluation of factual precision in long form text generation](#).
- Ramesh Nallapati, Bowen Zhou, Cicero Nogueira dos santos, Caglar Gulcehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence rnns and beyond](#).
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.

- Ani Nenkova, Rebecca Passonneau, and Kathleen McKeown. 2007. [The pyramid method: Incorporating human content selection variation in summarization evaluation](#). *ACM Trans. Speech Lang. Process.*, 4(2):4–es.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2019. [Adversarial NLI: A new benchmark for natural language understanding](#). *CoRR*, abs/1910.14599.
- Shon Otmazgin, Arie Cattan, and Yoav Goldberg. 2022. [F-coref: Fast, accurate and easy to use coreference resolution](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 48–56, Taipei, Taiwan. Association for Computational Linguistics.
- Alicia Parrish, William Huang, Omar Agha, Soo-Hwan Lee, Nikita Nangia, Alexia Warstadt, Karmanya Agarwal, Emily Allaway, Tal Linzen, and Samuel R. Bowman. 2021. [Does putting a linguist in the loop improve NLU data collection?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4886–4901, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *CoRR*, abs/1910.10683.
- Leonardo F. R. Ribeiro, Mengwen Liu, Iryna Gurevych, Markus Dreyer, and Mohit Bansal. 2022. [FactGraph: Evaluating factuality in summarization with semantic graph representations](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3238–3253, Seattle, United States. Association for Computational Linguistics.
- Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, Jacopo Staiano, Alex Wang, and Patrick Gallinari. 2021. [QuestEval: Summarization asks for fact-based evaluation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6594–6604, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alessandro Scirè, Simone Conia, Simone Ciciliano, and Roberto Navigli. 2023. [Echoes from alexandria: A large resource for multilingual book summarization](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 853–867, Toronto, Canada. Association for Computational Linguistics.
- Noam Shazeer and Mitchell Stern. 2018. [Adafactor: Adaptive learning rates with sublinear memory cost](#). *CoRR*, abs/1804.04235.
- Chenhui Shen, Liying Cheng, Xuan-Phi Nguyen, Yang You, and Lidong Bing. 2023. [Large language models are not yet human-level evaluators for abstractive summarization](#).
- Derek Tam, Anisha Mascarenhas, Shiyue Zhang, Sarah Kwan, Mohit Bansal, and Colin Raffel. 2022. [Evaluating the factual consistency of large language models through summarization](#).
- Liyan Tang, Tanya Goyal, Alex Fabbri, Philippe Laban, Jiacheng Xu, Semih Yavuz, Wojciech Kryscinski, Justin Rousseau, and Greg Durrett. 2023. [Understanding factual errors in summarization: Errors, summarizers, datasets, error detectors](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11626–11644, Toronto, Canada. Association for Computational Linguistics.
- Alex Wang, Kyunghyun Cho, and Mike Lewis. 2020. [Asking and answering questions to evaluate the factual consistency of summaries](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5008–5020, Online. Association for Computational Linguistics.
- Jiaan Wang, Yunlong Liang, Fandong Meng, Zengkui Sun, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. 2023a. [Is chatgpt a good nlg evaluator? a preliminary study](#).
- Yiming Wang, Zhuosheng Zhang, and Rui Wang. 2023b. [Element-aware summarization with large language models: Expert-aligned evaluation and chain-of-thought method](#).
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2017. [A broad-coverage challenge corpus for sentence understanding through inference](#). *CoRR*, abs/1704.05426.
- Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. [AlignScore: Evaluating factual consistency with a unified alignment function](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11328–11348, Toronto, Canada. Association for Computational Linguistics.
- Shiyue Zhang and Mohit Bansal. 2021. [Finding a balanced degree of automation for summary evaluation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6617–6632, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#).

A Prompt for claim extraction.

Figure 2 illustrates the formulated prompt designed for extracting claims from a given summary. This prompt was employed to generate the claims utilized in our evaluation metric and to train our distilled claim generation model.

B Claim extraction evaluation with BERTScore.

In Table 8, we report the claim extraction evaluation of the LLM-based claim extractor compared to its T5-based distilled version, using BERTScore instead of ROUGE.

Model	Precision (%)	Recall (%)	F1 (%)
GPT-3.5	95.8	94.3	95.0
T5 _{dist_GPT}	95.5	93.9	94.7

Table 8: Claim extractor performance comparison on ROSE (Liu et al., 2023b) using BERTScore (Zhang et al., 2020) as backbone metric in Equations 9 and 10 instead of ROUGE.

C NLI model ablation.

We evaluate the performance of FENICE utilizing RoBERTa⁵ (Liu et al., 2019) as the NLI model for our alignments (FENICE_{RoBERTa}), instead of DeBERTa (Table 9). The results reported are directly comparable to the ones presented in Table 4.

Metric	CNN	XSUM	AVG	Δ AVG
FENICE	67.5 \pm 3.0	73.7 \pm 0.3	70.6	-
FENICE _{RoBERTa}	64.4 \pm 3.2	68.1 \pm 1.6	66.3	-4.3

Table 9: Ablation study. We compare the performance of FENICE on the AGGREGFACT-FTSOTA split using DeBERTa and RoBERTa as the backbone NLI model.

D Long-form summarization factuality annotations statistics.

In Table 10, we report the number of factual and unfactual summaries in our collection.

Factual	Not Factual
66	273

Table 10: Number of factual/unfactual summaries in our set of human annotations.

⁵We load the pre-trained model weights from the Hugging Face Hub, available at: [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 define a claim as an "elementary information unit in a sentence, which no longer needs to be further split."

For example, given the following sentence:

INPUT:

NASA's Perseverance rover has discovered ancient microbial life on Mars according to a recent study published in the journal Science. It established a set of new paradigms for space exploration"

OUTPUT:

```
{'claims': ["NASA's Perseverance rover discovered ancient microbial life.", "Ancient microbial life was discovered on Mars.", "The discovery was made according to a recent study.", "The study was published in the journal Science.", "The study established a set of new paradigms for space exploration."]}
```

Please consider the following recommendations:

If possible, use a noun as the subject in the claim (try to avoid pronouns), such as "The study established a set of new paradigms for space exploration." in the previous example.

Do not generate any novel word, be faithful to the provided input.

Your response must be directly the json, do not add any other text, such as "Here is the output:" and similar. Note that, each fact expressed in the source text must be present in the output.

Now do this task for this input

INPUT:

Figure 2: LLM prompt for extracting claims given a summary.