

MEMERAG: A Multilingual End-to-End Meta-Evaluation Benchmark for Retrieval Augmented Generation

María Andrea Cruz Blandón^{1*}, Jayasimha Talur^{*} Bruno Charron,
Dong Liu, Saab Mansour, Marcello Federico
¹Tampere University, Amazon

Abstract

Automatic evaluation of retrieval augmented generation (RAG) systems relies on fine-grained dimensions like faithfulness and relevance, as judged by expert human annotators. Meta-evaluation benchmarks support the development of automatic evaluators that correlate well with human judgement. However, existing benchmarks predominantly focus on English or use translated data, which fails to capture cultural nuances. A native approach provides a better representation of the end user experience.

In this work, we develop a Multilingual End-to-end Meta-Evaluation RAG benchmark (MEMERAG). Our benchmark builds on the popular MIRACL dataset, using native-language questions and generating responses with diverse large language models (LLMs), which are then assessed by expert annotators for faithfulness and relevance. We describe our annotation process and show that it achieves high inter-annotator agreement. We then analyse the performance of the answer-generating LLMs across languages as per the human evaluators. Finally we apply the dataset to our main use-case which is to benchmark multilingual automatic evaluators (LLM-as-a-judge). We show that our benchmark can reliably identify improvements offered by advanced prompting techniques and LLMs¹.

1 Introduction

Retrieval augmented generation (RAG) is emerging as a popular application of large language models (LLMs) and a powerful paradigm to improve LLMs factuality (Gao et al., 2024). A RAG pipeline first retrieves relevant documents from an index based on a query and then composes a response

using an LLM. Grounding the LLM response on retrieved knowledge helps mitigate outdated knowledge, lack of domain expertise and reduce hallucinations (Lewis et al., 2020; Gao et al., 2024). Collecting benchmarking data for RAG is challenging due to the complexity of the pipeline that includes *information retrieval* and *text generation*. Text generation, our focus in this paper, has, in general, two modes of *automatic evaluation*: reference-based and reference-free, which differ in the availability of human-generated gold references for each model input. Both modes can either leverage single (e.g. BERTScore (Zhang et al., 2020)) or multidimensional (e.g. autoMQM (Fernandes et al., 2023)) scores. Multidimensional evaluation (Burchardt, 2013) provides more comprehensive understanding of text generation systems and is the de facto standard in the machine translation (MT) community².

In a reference-free evaluation setup, gold multidimensional judgements (factuality, relevance, etc) of model generations can be leveraged in a *meta-evaluation* framework. In this framework, automated evaluators are evaluated against human judgement to measure correlation. The automated evaluators can then be applied to measure the performance of new models' output.

Previous work for RAG meta-evaluation mainly focused on English (Fan et al., 2024) or leveraged human or machine translation of English datasets (Sharma et al., 2024). Multilingual meta-evaluation is important to reliably measure performance across languages which can vary depending on language characteristics (low vs. high resource, complex morphology, etc.) and scripts (Latin vs. non-Latin). Translation-based benchmarks, while permitting cross-language comparisons, suffer from translationese phenomena such as introducing simpler syntax and lexical choices

^{*}Work done by first author during an internship at Amazon. First two authors contributed equally to this work. Corresponding author: Jayasimha Talur email talurj@amazon.com.

¹We release <https://github.com/amazon-science/MEMERAG> our benchmark to support the community developing accurate evaluation methods for multilingual RAG systems.

²Since 2021 in the WMT metrics shared task <https://www2.statmt.org/wmt24/metrics-task.html>

(Baker et al., 1993; Graham et al., 2020), thus leading to data distributionally different from native data and not necessarily reflecting native users preferences (Chen et al., 2024). Our position is that translation-based (parallel) benchmarks should be complemented by native multilingual benchmarks.

To bridge those gaps, we propose a native meta-evaluation multilingual benchmark for RAG systems. Our benchmark is built on top of the popular MIRACL (Zhang et al., 2023) dataset³ that includes native questions across 18 languages and relevance judgements of retrieved passages for multilingual retrieval evaluation. We extend MIRACL by generating answers in five languages with a diverse set of LLMs, and collecting judgements on the faithfulness and relevance of the answers using native expert human annotators. For the latter, we devised a structured annotation process that achieved a high rate of inter-annotator agreement. To evaluate the benchmark and set reference baseline results for others to compare against, we run LLM-as-a-judge experiments with various prompting techniques and state-of-the-art LLMs.

To summarize, our main contributions are:

- We built and publicly release the first (to the best of our knowledge) native multilingual meta-evaluation RAG benchmark.
- We developed a rigorous flow chart-based annotation process to achieve high inter-annotator agreement rate for both faithfulness and relevance judgements.
- We evaluated the quality of the benchmark on three multi-lingual meta-evaluation aspects: prompt selection, model selection, and fine-grained analysis.
- We establish reference baselines of multilingual automatic evaluators on our benchmark, showcasing performance improvements when using advanced prompting and LLMs.

2 Related Work

Due to the pipeline approach of RAG systems, the evaluation can be split into 3 main components: 1) retriever metrics to identify relevant chunks of information to the input typically measured with recall/precision@K (Manning et al., 2008); 2) generator metrics to identify the “usefulness” of the generated answers in relation to the input, this is done across fine grained dimensions such as faithfulness and relevance, either leveraging references

answers (Es et al., 2024); 3) end-to-end (overall) metrics that take into account additional components such as preprocessing, chunking, query reformulation and cascading errors. Retrieval metrics have been extensively studied in the information retrieval community, hence recent work focused on the text generation performance of RAG systems. This is also the focus of our work. One important difference between generation with and without retrieved documents is the conflict between the parametric “world knowledge” and the non-parametric retrieved documents knowledge, hence the distinction between faithfulness against the retrieved documents (RAG specific) and factuality according to general knowledge (Maynez et al., 2020; Wu et al., 2024).

Recent studies have investigated the importance of various components within multilingual RAG systems. (Chirkova et al., 2024) utilized existing multilingual QA datasets to evaluate different combinations of retrievers and generator models, finding that task-specific prompt engineering is crucial for high-quality multilingual generation. In another study, (Thakur et al., 2024) extend MIRACL dataset to develop “MIRAGE-BENCH” a synthetic arena-based benchmark for ranking multilingual LLM generators. They employed preference judgments from a GPT-4o “judge” to train a ranking model. However, an important limitation of such synthetic benchmarks is the potential for self-preference bias (Panickssery et al., 2024), where the LLM judge may favor its own generations.

In this work, we focus on the faithfulness and relevance aspects to ensure meaningful results. These dimensions are typically assessed by human evaluators based on model-generated outputs, highlighting the need for developing automatic evaluation metrics that correlate well with human judgment—a process known as meta-evaluation. This need has led to a recent trend in the English-language research community of publishing meta-evaluation datasets and developing automated evaluators (Es et al., 2024; Saad-Falcon et al., 2024).

To the best of our knowledge, no multilingual meta-evaluation benchmark for RAG systems currently exists. In this work, we address this gap by developing such a benchmark. Our dataset facilitates the creation of multilingual automatic evaluators that correlate well with human judgments. This, in turn, enables comprehensive end-to-end benchmarking of RAG systems. We believe our dataset is the first to offer this capability in a multi-

³<https://huggingface.co/datasets/miracl/miracl>

MIRACL dataset Retrieval benchmark						
Lang.	Question	LLM response	Human judgement			
			Faithful		Relevant	
EN	When did Julia Child write her book Mastering the Art of French Cooking?	She co-authored it with Simone Beck and Louisette Bertholle.	✓	Logical conclusion	✓	Adds context to the answer
DE	Welche bekannte Universität gibt es in Pennsylvania?	Hier sind einige Beispiele:	X	Contradiction	X	Unrelated to the question
ES	¿Cómo la ciudad de Chicago tiñe de verde el río para celebrar San Patrick?	Ahora se utiliza un tinte vegetal seguro y no perjudicial para los organismos vivos en el río.	✓	Direct paraphrase	✓	Adds context to the answer
FR	Quels sont les trains français avec la vitesse la plus haute?	Les trains S-103 et CRH 3, tous deux de type Velaro, sont les trains avec la vitesse d'exploitation la plus élevée.	X	Adds new information	✓	Directly answers the question
HI	कैंसर के उपचार में प्रयुक्त उत्कृष्ट गैस कौनसी है?	दिए गए पाठान्तों में कैंसर के उपचार के लिए किसी गैस का उल्लेख नहीं है।	✓	Logical conclusion	✓	Directly answers the question
MEMERAG (ours) Meta-evaluation benchmark						

MEMERAG (ours) Meta-evaluation benchmark

Figure 1: Examples from our Multilingual End-to-end Meta-Evaluation for RAG (MEMERAG) dataset. We select MIRACL native multilingual questions (for a subset of 5 languages), generate responses using diverse LLMs, and annotate each generated sentence with human experts for faithfulness and relevance. The annotation includes both coarse-grained (✓, X) and fine-grained labels. Our dataset forms a meta-evaluation benchmark where automated evaluators can be developed and assessed on their correlation to human judgement. Due to space constraints we omit the retrieved documents (context) for the question and show only one sentence per LLM response.

lingual context.

3 Dataset Construction

Meta-evaluation datasets enable the development of reliable automatic evaluators. In a RAG setup, the input to the evaluator is composed of a question q (user input), a context c (set of passages automatically retrieved to the question) and an answer a generated by a language model to answer the question based on the context. The end-to-end evaluator then needs to judge the quality of the answer a given the context and the question (c, q) . Following previous work (Saad-Falcon et al., 2024; Es et al., 2024) we focus on two quality dimensions:

Faithfulness Is the answer grounded on the context, regardless of your world knowledge?

Relevance Is the answer relevant to the question, regardless of the context?

We build a multilingual end-to-end meta-evaluation RAG (MEMERAG) dataset by extending the MIRACL dataset (Zhang et al., 2023) to include *model-generated answers and human-based quality judgements*. More precisely, we select relevant question-context pairs, generate answers using various language models and gather expert human annotations on the quality of those an-

swers. Our dataset encompasses 5 languages: English (EN), German (DE), Spanish (ES), French (FR), and Hindi (HI), which represent multiple language families and both high- and low-resource languages. Figure 1 shows examples from the dataset, with LLM-generated answers and coarse- to fine-grained human-assigned labels for the faithfulness and relevance dimensions.

The MIRACL dataset is composed of questions written by humans in their *native* languages, one or more passages automatically retrieved from the Wikipedia, and human annotations about the relevance of each passage. Building a dataset starting from native questions in each language allows to evaluate RAG pipelines without resorting to (machine) translations, thus avoiding limitations and biases associated with translation. Note, however, that as questions were elicited from native speakers independently across different languages, the resulting data set is not parallel.

3.1 Question Selection

The questions in the MIRACL dataset were generated by humans based on prompts. This leads to questions that may be answerable by the prompt but may have ambiguity outside of that context. In particular, we identified as problematic the questions

for which the right answer can change over time. For example “*Who is the president of Spain?*” or “*How old is Drake Hogestyn?*”. In a RAG setting, different passages may have been written at different times and provide conflicting context. Additionally, the time of reference is usually not explicit in the question. To remove those complications we automatically filtered out time-dependent questions across all languages⁴. We combined the train and dev splits of the MIRACL dataset, which corresponds to a total of 3,662, 305, 2,810, 1,486 and 1,519 questions respectively for EN, DE, ES, FR and HI. Of those, 244, 13, 120, 64 and 86 were identified as time-dependent and filtered out.

3.2 Context Selection

The MIRACL dataset has an average of 10.3 passages per question, which corresponds to an average of 1,218 words of context in the English train and dev splits, with similar numbers in other languages. To reduce the cognitive load on human annotators, we limit the context per query to 5 passages.

The source dataset provides human-annotated binary relevance labels for passages. To more accurately simulate an automated retrieval process, we rank the passages for each question using BM25 (Schütze et al., 2008), as implemented in (Lü, 2024). We then select the top-5 ranked passages for each question. If these top-5 passages do not contain any human-annotated relevant passages, we replace the lowest-ranked passage with the highest-ranked relevant passage from the full set. This approach ensures that each question has at least one relevant passage in its context, avoiding scenarios where annotators would evaluate responses without any relevant information.

It is worth noting that simulating scenarios where no relevant passages exist is straightforward (e.g., by including only irrelevant passages). In such cases, for faithfulness evaluation, we would expect responses like “The provided documents do not contain a relevant answer.” Our method focuses on faithfulness while efficiently utilizing human annotation efforts by ensuring that each evaluated case has at least some relevant context.

3.3 Answer Generation

After question and passage selection, we generate an answer for each question-context pair and each

⁴See Appendix F for all prompts used during dataset construction.

of five state-of-the-art LLMs. We generated answers with Claude 3 Sonnet, Llama3 70B, Llama3 8B, Mistral 7B, and GPT-4o mini. Those LLMs were selected to cover a range of model sizes, open weight and proprietary models. We prompted all the models in English⁵, asking to answer the question based only on the given context, and requesting the answer to be provided in the same language as the context and question. For all models, we set the temperature to 0.1, and maximum number of output tokens to 1000.

We thus produced answers for more than 1000 questions per language, except for German for which MIRACL only contains 305 questions. As our focus is on long-form answers, we further filtered out questions for which any of the 5 models generated an answer shorter than 10 words.

3.4 Annotation Guidelines

The task of annotating answers with faithfulness is challenging due to several factors. First, it involves some subjectivity which might impact Inter-Annotator Agreement (IAA) (Kryscinski et al., 2020; Tang et al., 2024b). Then, its label space is not precisely defined in the literature (Tang et al., 2024b; Laban et al., 2023; Malaviya et al., 2024). Finally, although faithfulness should be ideally evaluated for atomic facts, it is generally evaluated at the sentence or even document level, due to annotation costs.

Starting with the factuality error taxonomy introduced in Tang et al. (2024b), we ran a number of annotation pilots to refine the label space and guidelines.

Finally, we converged to three coarse-grained labels (*Supported*, *Not supported*, *Challenging to determine*), explained through 10 fine-grained labels. To increase the consistency of the annotation (IAA), we guide the annotation process through a flow chart (documented in Figure 3 in Appendix A). For relevance, which is significantly less ambiguous to evaluate, we devise a simple annotation process with three labels: *Directly answers the question*, *Adds context to the answer*, and *Unrelated to the question*. Note that the first two labels can be used to describe “relevant” sentences, while the last label identifies “irrelevant” sentences. (See Appendix A

⁵There is evidence in the literature for better model accuracy when the models are prompted to process in English (Lai et al., 2023; Liu et al., 2024), though under particular scenarios other languages could perform better (Behzad et al., 2024). We leave additional prompting experiments for future work.

Lang	#Q	Answer		Context
		#S	Avg. #W	Avg. #W
EN	250	400	30.3	613.5
DE	250	468	27.3	455.0
ES	250	563	52.1	522.3
FR	250	540	48.7	478.3
HI	250	351	23.8	571.5
Total	1,250	2,322		

Table 1: General statistics of the MEMERAG dataset. #Q: number of questions, #S: number of sentences, #W: number of words. Each answer is annotated at the sentence level, leading to 2,322 total sentences annotated by experts for faithfulness and relevance.

for more details.)

3.5 Annotation Process

From the question-context-answer triplets obtained in Section 3.3, we randomly sampled 250 questions per language (50 per answer-generating model, without overlapping questions for diversity). We employed a professional vendor with native annotators⁶ to gather annotations for each sentence⁷ of the 250 answers per language. The statistics of the annotated dataset are presented in Table 1. Among the 250 answers per language, a random subset of 10 were assigned to 3 annotators for computing the IAA and the rest to a single annotator (see Table 2).

The annotations were gathered via a web-based tool that implemented the flow chart of the annotation guidelines (see Appendix A for details). To further enhance IAA, we drew upon the findings of Krishna et al. (2023), which demonstrated that highlighting relevant information aids annotators in performing tasks and reaching consensus. Thus, we utilized the Llama 3 70B LLM to identify sentences within the retrieved passages that could potentially serve as supporting information to the answer sentences (see Appendix F).

The English annotations required approximately 25 hours of total annotation time, averaging 5.5 minutes per question. This covered 250 questions, including 10 that were annotated by three different annotators for quality control. Similar time investments were observed for the other four languages.

Table 2 summarizes the IAA per language for faithfulness and relevance labels assigned by 3 an-

⁶Annotators were compensated with a competitive hourly rate that is benchmarked against similar roles in their country of residence.

⁷Sentences were segmented using the pySBD(Sadvilkar and Neumann, 2020) package.

notators. We report IAA using Gwet’s AC1 (Gwet, 2008) and Fleiss Kappa (Fleiss, 1971). We observe high agreement for faithfulness (0.84-0.93 Gwet’s AC1 and 0.70-0.88 Fleiss Kappa) and even higher agreement for relevance (0.95-1.0 Gwet’s AC1 and 0.63-1.0 Fleiss Kappa). This shows that the annotators are aligned and indicates a high quality of the annotations. Comparing to previous work, Tang et al. (2024b) report a Fleiss Kappa of 0.34-0.42 on faithfulness labels which they deem fair to moderate agreement. Note that a direct comparison with this work is not possible as they deal with different tasks, nevertheless the high IAA we are reporting is a testament of the effectiveness of our flow chart-based annotation design. Further details on IAA with fine-grained explanatory labels, and an extended version of the dataset (MEMERAG-Ext) with five annotators for 150 answers per language can be found in Appendix B.

Lang	Faithfulness		Relevance	
	Gwet’s AC1	Fleiss Kappa	Gwet’s AC1	Fleiss Kappa
EN	0.93	0.77	1.00	1.00
DE	0.84	0.81	0.95	0.73
ES	0.91	0.76	1.00	1.00
FR	0.89	0.88	0.93	0.63
HI	0.89	0.70	1.00	1.00

Table 2: Inter-annotator agreement (IAA) on the faithfulness and relevance dimensions with 3 annotators. Annotations are at the answer sentence level.

4 Annotation Results

We present in this section the results of the human annotations for the 2,322 sentences of the MEMERAG dataset.

Table 3 shows the distribution of faithfulness and relevance labels across the five languages. The distribution of labels is consistent across all languages, with a few exceptions. On faithfulness, German and Hindi show higher percentages of *Supported* answers. On relevance, Spanish presents a significantly higher percentage of labels *Adds context to the answer* compared to other languages while English had a very low share of labels *Unrelated to the question*. As a partial explanation for this, we note that Spanish questions generated the largest numbers of output sentences with 52.1 words vs., for example, 30.3 words for English (see Table 1).

Lang	Faithfulness			Relevance		
	✓	×	?	✓	✓	×
EN	65.2	31.5	3.2	65.2	32.5	2.2
DE	71.2	26.7	2.1	61.3	26.5	12.2
ES	65.7	32.9	1.4	48.8	43.9	7.3
FR	62.0	37.8	0.2	63.3	29.3	7.4
HI	73.8	25.6	0.6	68.9	21.4	9.7

Table 3: Label distribution in the benchmark. Percentage of sentences labelled as supported (✓), not supported (×), challenging to determine (?) for faithfulness, and as directly answers the question (✓), adds context to the answer (✓), unrelated to the question (×) for relevance, per language.

Label	EN	DE	ES	FR	HI
Direct paraphrase	8.5	41.7	25.0	33.7	40.7
Logical conclusion	41.2	28.2	30.4	28.0	5.7
Other	15.5	1.3	10.3	0.4	27.4
Adds new info	7.0	9.6	16.0	15.0	14.8
Contradiction	4.5	11.3	8.3	5.9	7.1
Mis-referencing	1.5	3.0	2.3	3.7	0.3
Nuance shift	6.8	0.6	4.3	5.6	1.7
Opinion as fact	0.5	0.6	0.2	2.2	0.3
Wrong reasoning	10.0	0.6	1.4	1.9	0.3
Other	1.2	0.9	0.4	3.5	1.1
Challeng. to determ.	3.2	2.1	1.4	0.2	0.6

Table 4: Fine-grained faithfulness label distribution in the benchmark. Percentage of sentences with each label per language. The three sections correspond to the coarse-grained labels *Supported/Not supported/Challenging to determine*.

Hence, we expect that the relevance statistics reflect the tendency of Spanish answers to be more verbose.

For a more granular insights into the annotations, we show in Table 4 the distribution of the explanatory fine-grained faithfulness labels for each language. As label names are quite self-explanatory, we refer for their precise meaning to Figure 3 in Appendix A. Table 4 is split into three blocks, respectively, addressing fine-grained labels for *Supported* (top), *Not supported* (middle), and *Challenging to determine* (bottom) answers. We observe remarkable differences across languages. For example, supported answers for English are prominently under form of a *logical conclusion* from the context (41.2%), for German and Hindi from *direct paraphrasing* of information in the context (41.7% and 40.7%), while for French and Spanish from a more balanced combination of the two reasons. On the

side of unsupported answers, the main mistake type in English is *Wrong reasoning* (10%), i.e. answers are non logical conclusions from the context, while this type of error is significantly rarer (under 2%) for all other languages. The rate of *Adds new information* errors, a.k.a. hallucinations, ranges from 7% for English to 16% for Spanish. The observed cross-linguistic variations in the distribution of labels can be attributed to the different nature of the questions and accuracy of the models across the 5 languages.

5 Dataset Applications

The MEMERAG dataset is designed to support the development of reliable automatic evaluation methods. For that purpose, we describe in this section how our dataset can be used as a benchmark to enable various meta-evaluation use cases. We focus on two applications: 1) **Prompt selection**: The ability of our benchmark to effectively select prompts for automatic evaluation, 2) **Model selection**: The effectiveness of our benchmark to compare and select models. By concentrating on these aspects, we can evaluate the benchmark’s utility as a comprehensive tool for multilingual model assessment. While we provide reference baselines for each application, our focus is on showcasing the effectiveness of the benchmark rather than the underlying capabilities of the LLMs.

5.1 Experimental Setup

Benchmark Tasks Our benchmark is composed of multiple tasks defined by the annotation dimension and subset considered. On the annotation dimension, we focus our experiments on the coarse-grained faithfulness dimension. This dimension is more challenging than relevance as highlighted by the lower IAA, while retaining a high-enough IAA to make for a trustworthy benchmark. We invite benchmark users to also experiment on the other dimensions provided by the dataset depending on their use case. Note that for this experiment, we remove the sentences labelled as *Challenging to determine* by human annotators. On the annotation subset, we first distinguish the *multilingual* task which uses the full dataset and the *monolingual* task, which only considers a single language. Those tasks can then be further broken down at the *fine-grained* level by considering the subset of sentences with a certain fine-grained label. Performance is evaluated with Balanced Accuracy

(BAcc), with equal weights on each coarse-grained label and language. We conduct significance testing using permutation tests (Good, 2013), with further details in Appendix E.

Reference Prompts To demonstrate how the benchmark can be used to select the appropriate prompt, we experiment with multiple prompting strategies from simple to advanced techniques, starting with zero-shot prompting, where LLMs directly classify statements as *Supported* or *Not supported*. We then implement chain-of-thought (COT) prompting (Wei et al., 2022), which incorporates an intermediate reasoning step. While these basic prompting strategies provide a good starting point for evaluation, our initial experiments revealed limitations in their ability to capture the nuanced requirements of our specific task. Without explicit guidelines in the prompt, automatic evaluators rely on their “world knowledge”, which may not align with the specific requirements of the evaluation task. To overcome this, we add instructions from the annotation guidelines (AG) in the prompt. Adding annotation guidelines provides clear criteria for what constitutes *Supported* versus *Not Supported* sentences, reducing ambiguity in the evaluation process. The various prompts are presented in Appendix G.

Reference Models We experiment with four LLMs with varying model sizes and capabilities: GPT-4o mini, Qwen 2.5 32B, and two versions of Llama 3.2 (11B and 90B)⁸. In case an LLM does not produce one of the required labels, we repeatedly prompt the LLM up to five times with temperature and top_p equal to 0.1 to get a valid label. If the LLM fails to generate a label after five retries, we treat the datapoint as an error (wrong label).

5.2 Experimental Results

Our benchmark enables systematic evaluation of different approaches to automated faithfulness evaluation. To illustrate this, we examine how the benchmark can surface the effectiveness of various automatic evaluation models and prompting strategies.

⁸In comparison to the LLMs selected for answer generation, see Section 3.3, we upgraded Llama from 3 to 3.2 as the context length of 8K tokens was not sufficient for all prompts. We picked GPT-4o mini as representative of proprietary models. Additionally, Mistral was excluded as it performed poorly in initial experiments.

Prompt	GPT-4o mini	Qwen 2.5 32B	Llama 3.2 90B	Llama 3.2 11B
ZS	59.7	66.7	58.0	55.4
COT	61.4	68.8	59.9	62.5
AG	71.6 [†]	72.6	62.8 [†]	57.9
AG+COT	71.7	71.8 [†]	64.4	61.6 [†]

Table 5: Reference baselines for the multilingual task on coarse-grained faithfulness. Balanced accuracy (BAcc) of the automatic evaluators averaged across the 5 languages (EN, DE, ES, FR, HI) using zero-shot (ZS), chain-of-thought (COT), annotation guidelines (AG) and AG+COT prompting strategies. Bold indicates best performance for the column, [†] indicates results not statistically different from the best ($p > 0.05$). Additional results on monolingual tasks and standard errors can be found in Appendix E, Tables 10-12.

Table 5 demonstrates the benchmark’s ability to compare different prompting approaches across languages. The benchmark reveals consistent patterns, showing how different prompt designs impact evaluation quality. As expected, adding a reasoning step (COT) improves over zero-shot prompting. In addition, adding annotation guidelines (AG) helps align automated evaluators with human judgments across all languages. Comparing the two best models GPT-4o mini and Qwen 2.5 32B, Qwen 2.5 32B excels in the zero-shot and COT setups, showcasing higher “out-of-the-box” alignment with human judgements. GPT-4o mini achieves similar performance once the annotation guidelines are added to the prompt.

Figure 2, shows the performance per language of various automatic evaluators with a fixed prompt (AG + COT), which allows us to select the best model for each language. We observe that GPT-4o mini performs best in English. For the rest of the languages Qwen 2.5 32B performs the best however, the results are not statistically different from GPT-4o mini. Our benchmark also provides users with the capability to conduct detailed, fine-grained analyses of model performance across various dimensions of faithfulness. The breakdown of automatic evaluation performance by error type is shown in Appendix I.

6 Conclusions

We introduced a high-quality and challenging multilingual end-to-end meta-evaluation benchmark for RAG (MEMERAG). Our carefully designed flow-chart-based annotation achieved a high inter-

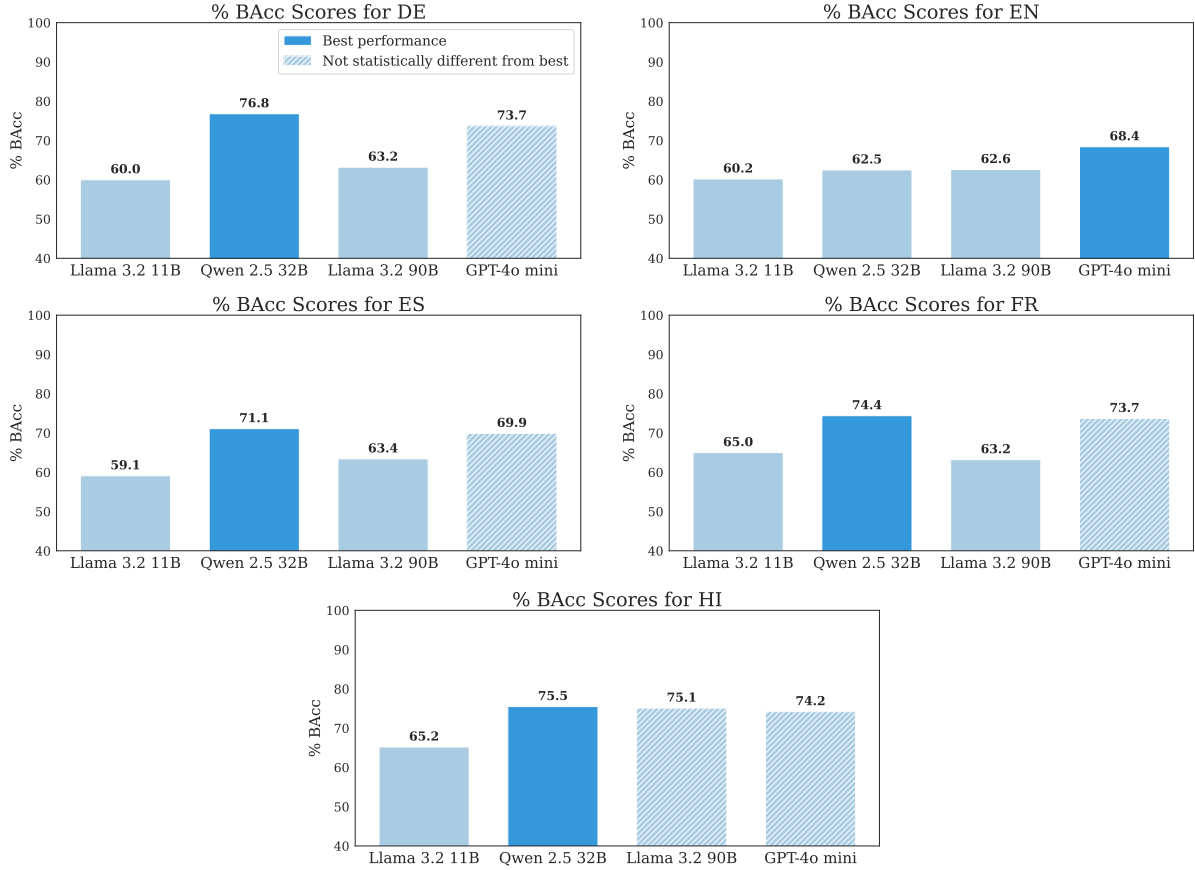


Figure 2: Reference baselines for the monolingual task on coarse-grained faithfulness. Balanced Accuracy (BAcc) of the automatic evaluators using various LLMs across five languages: EN, DE, ES, FR and HI. Each plot compares the performance of four models: Llama 3.2 11B, Qwen 2.5 32B, Llama 3.2 90B, and GPT-4o mini using AG + COT prompt. The best-performing model for each language is highlighted with a darker blue bar. Bars with diagonal hatching indicate results not statistically different from the best ($p > 0.05$).

annotator agreement rate supporting the reliability of the benchmark. The introduced MEMERAG dataset opens the door for multiple application scenarios, including but not limited to the demonstrated cases, i.e. prompt selection and model selection.

For the meta-evaluation setup, we develop and compare various LLMs-as-a-judge and observe that automatic evaluators performance varies across the languages, influenced by language characteristics, native-question complexity and LLM generation nuances. These variations underscore the importance of our testbed, which demonstrated consistent results when comparing the prompting techniques (COT+guidelines > COT > zero-shot) and provides a foundation for developing better multilingual evaluators.

7 Limitations

Due to time and cost constraints, our annotations and experiments are limited in terms of prompting

techniques, LLMs we experimented with and languages we annotated. Nevertheless, we diversified our LLMs across size and “openness” while the languages represent two families and low and high resource ones. In addition, there exists in the literature fine-tuned factuality evaluators for English, though we expect those to not work as well on non-English languages. Another method is to approximate factuality through entailment tasks (i.e. XNLI dataset) though such methods were shown (for English) to be inferior to multi-task training and distillation and data augmentation from LLMs (Tang et al., 2024a). Fine-tuning multilingual evaluators and examining transfer learning across languages is interesting but is left for future work that can leverage our dataset for this purpose.

As we advocate for a native testing approach, the questions across the languages are not parallel, which could introduce a dimension of different questions and LLM generations complexities across the different language test data. The data

we collected presents different challenges which are captured according to our fine grained error labels (Table 4). Future work could balance the challenges and complexities by collecting data for specific challenging phenomena. Note that this balancing is not straightforward, it can be done on the question side though this is insufficient as it does not control for the answer complexity. Controlling for the answer complexity is a challenging problem as the answer side is model generated (one method is to generate many answers and select for certain phenomena with human in the loop which is costly).

References

- Mona Baker, Gill Francis, and Elena Tognini-Bonelli. 1993. [Corpus linguistics and translation studies: Implications and applications](#). In *In Text and Technology: In Honour of John Sinclair*, Netherlands. John Benjamins Publishing Company.
- Shabnam Behzad, Amir Zeldes, and Nathan Schneider. 2024. [To ask LLMs about English grammaticality, prompt them in a different language](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15622–15634, Miami, Florida, USA. Association for Computational Linguistics.
- Aljoscha Burchardt. 2013. [Multidimensional quality metrics: a flexible system for assessing translation quality](#). In *Proceedings of Translating and the Computer 35*, London, UK. Aslib.
- Pinzhen Chen, Simon Yu, Zhicheng Guo, and Barry Haddow. 2024. [Is it good data for multilingual instruction tuning or just bad multilingual evaluation for large language models?](#) In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9706–9726, Miami, Florida, USA. Association for Computational Linguistics.
- Nadezhda Chirkova, David Rau, Hervé Déjean, Thibault Formal, Stéphane Clinchant, and Vassilina Nikoulina. 2024. [Retrieval-augmented generation in multilingual settings](#). *Preprint*, arXiv:2407.01463.
- Shahul Es, Jithin James, Luis Espinosa Anke, and Steven Schockaert. 2024. [RAGAs: Automated evaluation of retrieval augmented generation](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 150–158, St. Julians, Malta. Association for Computational Linguistics.
- Wenqi Fan, Yajuan Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and Qing Li. 2024. [A survey on rag meeting llms: Towards retrieval-augmented large language models](#). In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '24*, page 6491–6501, New York, NY, USA. Association for Computing Machinery.
- Patrick Fernandes, Daniel Deutsch, Mara Finkelstein, Parker Riley, André F. T. Martins, Graham Neubig, Ankush Garg, Jonathan H. Clark, Markus Freitag, and Orhan Firat. 2023. [The Devil is in the Errors: Leveraging Large Language Models for Fine-grained Machine Translation Evaluation](#). *arXiv preprint*. ArXiv:2308.07286 [cs].
- Joseph Fleiss. 1971. [Measuring nominal scale agreement among many raters](#). *Psychological Bulletin*, 76:378–.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. [Retrieval-Augmented Generation for Large Language Models: A Survey](#). *arXiv preprint*. ArXiv:2312.10997 [cs].
- Phillip Good. 2013. *Permutation tests: a practical guide to resampling methods for testing hypotheses*. Springer Science & Business Media.
- Yvette Graham, Barry Haddow, and Philipp Koehn. 2020. [Statistical power and translationese in machine translation evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 72–81, Online. Association for Computational Linguistics.
- Kilem Li Gwet. 2008. Computing inter-rater reliability and its variance in the presence of high agreement. *British Journal of Mathematical and Statistical Psychology*, 61(1):29–48.
- 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.
- Philippe Laban, Wojciech Kryscinski, Divyansh Agarwal, Alexander Fabbri, Caiming Xiong, Shafiq Joty, and Chien-Sheng Wu. 2023. [SummEdits: Measuring LLM Ability at Factual Reasoning Through The Lens of Summarization](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9662–9676, Singapore. Association for Computational Linguistics.
- Viet Dac Lai, Nghia Ngo, Amir Pouran Ben Veyseh, Hieu Man, Franck Dernoncourt, Trung Bui, and

- Thien Huu Nguyen. 2023. [ChatGPT beyond English: Towards a comprehensive evaluation of large language models in multilingual learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13171–13189, Singapore. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Chaoqun Liu, Wenxuan Zhang, Yiran Zhao, Anh Tuan Luu, and Lidong Bing. 2024. [Is translation all you need? a study on solving multilingual tasks with large language models](#). *Preprint*, arXiv:2403.10258.
- Xing Han Lù. 2024. [Bm25s: Orders of magnitude faster lexical search via eager sparse scoring](#). *Preprint*, arXiv:2407.03618.
- Chaitanya Malaviya, Subin Lee, Sihao Chen, Elizabeth Sieber, Mark Yatskar, and Dan Roth. 2024. [ExpertQA: Expert-Curated Questions and Attributed Answers](#). *arXiv preprint*. ArXiv:2309.07852 [cs].
- Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press, Cambridge, UK.
- 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.
- Arjun Panickssery, Samuel R. Bowman, and Shi Feng. 2024. [Llm evaluators recognize and favor their own generations](#). *Preprint*, arXiv:2404.13076.
- Jon Saad-Falcon, Omar Khattab, Christopher Potts, and Matei Zaharia. 2024. [ARES: An Automated Evaluation Framework for Retrieval-Augmented Generation Systems](#). *arXiv preprint*. ArXiv:2311.09476 [cs].
- Nipun Sadvilkar and Mark Neumann. 2020. [PySBD: Pragmatic sentence boundary disambiguation](#). In *Proceedings of Second Workshop for NLP Open Source Software (NLP-OSS)*, pages 110–114, Online. Association for Computational Linguistics.
- Hinrich Schütze, Christopher D Manning, and Prabhakar Raghavan. 2008. *Introduction to information retrieval*, volume 39. Cambridge University Press Cambridge.
- Nikhil Sharma, Kenton Murray, and Ziang Xiao. 2024. [Faux polyglot: A study on information disparity in multilingual large language models](#). *Preprint*, arXiv:2407.05502.
- Liyan Tang, Philippe Laban, and Greg Durrett. 2024a. [MiniCheck: Efficient Fact-Checking of LLMs on Grounding Documents](#). *arXiv preprint*. ArXiv:2404.10774 [cs].
- Liyan Tang, Igor Shalyminov, Amy Wing-mei Wong, Jon Burnsky, Jake W. Vincent, Yu’an Yang, Siffi Singh, Song Feng, Hwanjun Song, Hang Su, Lijia Sun, Yi Zhang, Saab Mansour, and Kathleen McKeown. 2024b. [TofuEval: Evaluating Hallucinations of LLMs on Topic-Focused Dialogue Summarization](#). *arXiv preprint*. ArXiv:2402.13249 [cs].
- Nandan Thakur, Suleman Kazi, Ge Luo, Jimmy Lin, and Amin Ahmad. 2024. [Mirage-bench: Automatic multilingual benchmark arena for retrieval-augmented generation systems](#). *Preprint*, arXiv:2410.13716.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed H. Chi, Quoc Le, and Denny Zhou. 2022. [Chain of thought prompting elicits reasoning in large language models](#). *CoRR*, abs/2201.11903.
- Kevin Wu, Eric Wu, and James Zou. 2024. [Clasheval: Quantifying the tug-of-war between an LLM’s internal prior and external evidence](#). In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Xinyu Zhang, Nandan Thakur, Odunayo Ogundepo, Ehsan Kamalloo, David Alfonso-Hermelo, Xiaoguang Li, Qun Liu, Mehdi Rezagholizadeh, and Jimmy Lin. 2023. [MIRACL : A Multilingual Retrieval Dataset Covering 18 Diverse Languages](#). *Transactions of the Association for Computational Linguistics*, 11:1114–1131.

A Human Annotation Guidance

Before conducting large-scale annotations, we conducted a pilot with 10 English RAG outputs and 3 annotators. We asked annotators to evaluate faithfulness at the sentence level either as *Supported* or with a subset of the factuality mistakes typology in Tang et al. (2024b) (developed for summarization): *Contradiction*, *Hallucination*, *Misreferencing*, *Nuance meaning shift*, *Opinion stated as fact*, *Wrong reasoning*, to which was an *Other mistake* label was added to account for unforeseen mistakes in the RAG setting. As this led to very low IAA, we conducted another round reducing the non-supported labels to *Stating opinion as fact*, *Drawing wrong conclusions*, *Other mistake* but this still resulted in low IAA (Gwet’s AC1 0.45).

Upon careful analysis of the annotator disagreements in the pilot and further rounds of calibration, we developed the annotation workflow shown in Figure 3. The key improvements were: (i) add a *Challenging to determine* label, (ii) have two levels of labels, a coarse-grained level (*Supported*, *Not supported*, *Challenging to determine*) and a fine-grained level for precision on the mistakes, (iii) enforce annotators to follow a specific reasoning with a flow chart, (iv) use numbers for the fine-grained level rather than labels which could be misinterpreted. Those guidelines allowed to reach significantly higher IAA on faithfulness (Gwet’s AC1 0.81 coarse-grained). Likewise we iterated on the relevance labels, starting from *Must have*, *Nice to have*, or *Irrelevant* and converging to the more explicit *Directly answers the question*, *Adds context to the answer*, or *Unrelated to the question*. This also increased the IAA significantly. The screenshot of the user interface used by human annotators is shown in Figure 4.

B IAA and Extended Dataset

As explained in Section 3.5, the IAA on the MEMERAG dataset was measured by triple annotation on a random subset of 10 questions per language. Detailed IAA metrics, at both coarse- and fine-grained levels, can be found in Table 6. We observe that IAA is high for faithfulness at coarse-grained level, and for relevance at both levels. This prompted us to 1) focus the experiments described in this paper to those dimensions with high IAA, and 2) investigate the low IAA for fine-grained faithfulness.

The disagreements highlighted by low IAA are

MEMERAG					
Lang	#S	IAA (Gwet’s AC1)			
		Faithfulness		Relevance	
		Coarse	Fine	Coarse	Fine
EN	13	0.93	0.63	1	0.92
DE	31	0.84	0.47	0.95	0.91
ES	20	0.91	0.3	1	0.93
FR	23	0.89	0.58	0.93	0.92
HI	17	0.89	0.18	1	1

Table 6: Inter-annotator agreement (IAA) on the faithfulness and relevance dimensions with 3 annotators (MEMERAG subset), for coarse-grained and fine-grained levels. Annotations are at the answer sentence level (#S number of answer sentences is provided).

MEMERAG-Ext					
Lang	#S	IAA (Gwet’s AC1)			
		Faithfulness		Relevance	
		Coarse	Fine	Coarse	Fine
EN	226	0.83	0.41	1.00	0.9
DE	272	0.75	0.47	0.92	0.79
ES	276	0.91	0.39	1.00	0.97
FR	370	0.72	0.53	0.99	0.8
HI	208	0.91	0.33	0.98	0.92

Table 7: Inter-annotator agreement (IAA) on the faithfulness and relevance dimensions with 5 annotators (MEMERAG-Ext), for coarse-grained and fine-grained levels. Annotations are at the answer sentence level (#S number of answer sentences is provided).

due to ambiguity in the annotation guidelines. The fact that such ambiguity remained after significant work on refining the guidelines and guiding annotators with UI elements hints at the irreducible subjectivity of the task. Subjectivity manifests itself as a non-trivial probability distribution of the annotation label obtained by drawing a random human annotator for a given question. It would be of interest in the development of automated annotators to compare this distribution to that of the annotations of a stochastic LLM-based annotator.

To investigate those aspects, we created an extended dataset, called MEMERAG-Ext. The dataset is composed of 150 question-context pairs per language, extracted similarly to MEMERAG with the process described in Section 3. The questions are disjoint between the two datasets, except for German where there is some overlap due to the limited number of questions in MIRACL. The annotation process then differed in two ways: 1) annotations were obtained for all question from 5

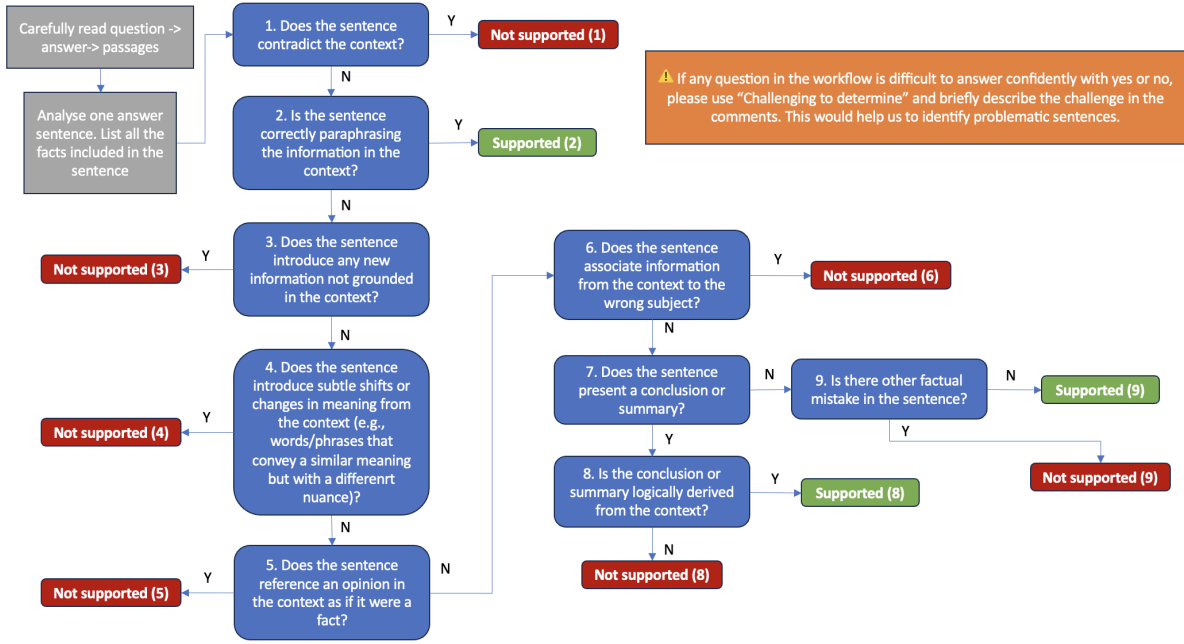


Figure 3: Annotation guideline for faithfulness labelling by human

Figure 4: User interface used by human annotators for labelling faithfulness and relevance.

expert annotators, and 2) The *Challenging to determine* label was not an option for annotators to force meaningful labels even at the cost of disagreements. IAA metrics for the MEMERAG-Ext dataset can be found in Table 7. Those results are in line with the IAA observed in MEMERAG, which confirms its quality. Detailed investigation of the disagreements between human annotators, and comparison to automated annotators are left as future work.

We release this extended dataset along with the main dataset to promote related work by the research community.

C Detailed numbers on the MEMERAG dataset

In this section, we look at the detailed statics of faithfulness and relevance in the MEMERAG

Lang	#Q	Answer		Context
		#S	Avg. #W	Avg. #W
EN	150	226	27.5	591.4
DE	150	272	30.0	456.8
ES	150	276	39.1	490.8
FR	150	370	48.1	453.2
HI	150	208	23.0	576.6
Total	750	1352		

Table 8: General statistics of the MEMERAG-Ext dataset. #Q: number of questions, #S: number of sentences, #W: number of words. Each answer is annotated at the sentence level, leading to 1,352 total sentences annotated by experts for faithfulness and relevance.

dataset itself. Overall, the faithfulness and relevance in the dataset variate a lot from one language to another language and from one generator model to another. There is usually a range of 10-20 percentage points between the language with lowest percentage and the language with highest percentage of supported sentences. This supports our assumption of variable behaviour/performance across different languages.

Given the large variation, the percentage of supported answers is larger than that of unsupported answers. Similarly, the ratio of relevant answers is in general larger than that of other relevance buckets (except for ES with LLM-D; EN, ES and FR with LLM-A). Interestingly, when analysing the various results per model and language, we did not find a general pattern. This is counter-intuitive to our assumption of having more supported sentences for English language. Surprisingly, even though Hindi is a low-resource language the models are producing 70% supported sentences, except for LLM-D that produced only around 62% supported sentences.

D Fine grained automatic annotation analysis

Figure 5 shows a heatmap representing the failure modes of automatic evaluators. We observe that improving the prompting strategy from ZS to AG + COT, reduces automatic evaluation errors across all error categories, except for Llama 3.2 11B, where the model makes more errors in the “Logical conclusion category”. We also observe that “Wrong reasoning”, “Nuance shift” and “Logical conclusion” are the top error categories for all the models tested. Future work could explore prompting or fine-tuning techniques designed to handle specific

error types.

E Statistical Significance Test Details

In Table 5 and Figure 2, we aim to determine whether the scores (in terms of BAcc) for the best prompt and best model were significantly different from the other scores in the table. To test statistical significance, we use permutation test (Good, 2013), a non-parametric method for comparing two related samples. The null hypothesis for this test suggests that there is no significant difference between the performance of the best-performing prompt/models and the other prompts/models, while the alternative hypothesis suggests a significant difference exists. We consider $\alpha = 0.05$ as significance level. Consequently, when $p > 0.05$, we are not able to reject the null hypothesis, indicating that prompts and models have similar performance.

F Prompts used for dataset construction

In this section we describe the various prompts used in our pipeline. Our prompts are written as Jinja2⁹ templates.

F.1 Time-dependent answer filtering

During internal pilots, we identified answers that relied on current date / current affairs that could be challenging to determine their faithfulness. For example, to the question "How old is Yann LeCun?" is the answer correct if the LLM uses the time when it was trained? To avoid such cases we used Llama3 70B to filter out those cases using the prompt shown below:

Time-dependent Answer filtering

You are an NLP assistant that helps to identify if a question requires to know when is today (day, or month, or year), current affairs, or up-to-date information. Give the step by step of how to answer the question in between the labels <rationale></rationale>. Then verify if the steps included to know any information about the current time and give your answer in between the tags: <label></label>. Please only use 'yes' or 'no' in your final answer. You will be given the question in {{language}}. Provide your rationale and label in English.

⁹<https://jinja.palletsprojects.com/en/stable/>

Model	EN	DE	ES	FR	HI
LLM-A (615)					
<i>Faithfulness</i>	52.7 / 43.8 / 3.5	74.3 / 17.1 / 8.6	58.6 / 41.4 / 0	59.6 / 40.4 / 0	73.2 / 26.8 / 0
<i>Relevance</i>	42.5 / 51.7 / 5.8	60.9 / 37.4 / 1.7	42.3 / 54.5 / 3.2	48 / 50.3 / 1.7	46.8 / 41.3 / 11.9
LLM-B (315)					
<i>Faithfulness</i>	76.1 / 17.9 / 6	88.4 / 11.6 / 0	88.2 / 11.8 / 0	58.6 / 41.4 / 0	84.9 / 15.1 / 0
<i>Relevance</i>	67.1 / 31.5 / 1.4	72.7 / 22.1 / 5.2	75 / 22.2 / 2.8	87.1 / 12.9 / 0	80.7 / 10.5 / 8.8
LLM-C (315)					
<i>Faithfulness</i>	61.4 / 35.1 / 3.5	69.8 / 30.2 / 0	68.1 / 31.9 / 0	80.3 / 19.7 / 0	80 / 20 / 0
<i>Relevance</i>	86.9 / 11.5 / 1.6	78.3 / 20.3 / 1.4	69.3 / 28 / 2.7	87.1 / 10 / 2.9	85.9 / 9.4 / 4.7
LLM-D (743)					
<i>Faithfulness</i>	67.3 / 30.8 / 1.9	61.4 / 38 / 0.6	61 / 35.7 / 3.3	59.8 / 40.2 / 0	63 / 34.6 / 2.5
<i>Relevance</i>	70.3 / 29.7 / 0	41.5 / 33 / 25.5	33 / 52.3 / 14.7	54.8 / 26.1 / 19.1	57.1 / 24.2 / 18.7
LLM-E (334)					
<i>Faithfulness</i>	77.2 / 21.1 / 1.7	73.9 / 26.1 / 0	69.7 / 29 / 1.3	59.2 / 39.5 / 1.3	73.3 / 26.7 / 0
<i>Relevance</i>	80.3 / 19.7 / 0	75.4 / 17.4 / 7.3	63.4 / 36.6 / 0	73.8 / 25 / 1.2	90.6 / 9.4 / 0

Table 9: Percentage of sentences labelled as *Supported/Not Supported/Challenging to determine* for the faithfulness, and as *Directly answers the question/Adds context to the answer/Unrelated to the question* for the relevance, per language and model. The total number of sentences generated by the model across all languages is given in parenthesis besides the model name.

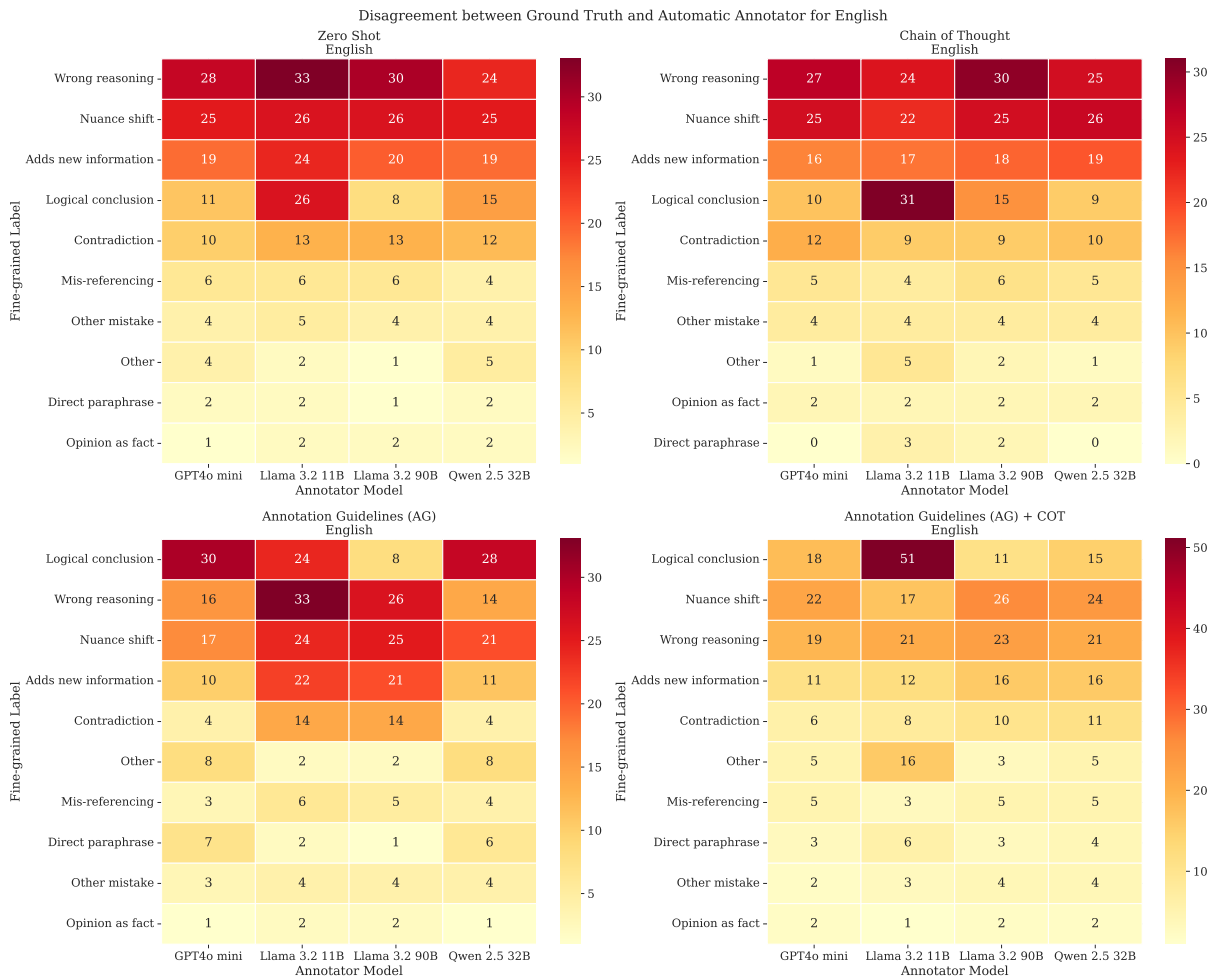


Figure 5: Fine-grained faithfulness errors when automatic evaluator disagrees with ground truth.

F.2 Highlighting relevant segments

We highlight relevant statements in the passage to help annotators focus on important parts of the context. We prompt Llama 3 70B, to predict all the statements that support the passage. We use temperature equal to 0.1 in this step.

F.3 Answer generation prompt

We generate the answers using the same prompt across all models and languages. All the instructions were given in English, while the context and question were given in the testing language. For all models we set the temperature to 0.1 and the maximum number of token to 1000.

Answer Generation Prompt

You are an NLP assistant whose purpose is to answer a question based on given passages. The passages may or may not help answer the question. You will need to provide the answer based only on the passages. The answer must be in language without fail. Be concise and direct without referring to passages in the answer. Avoid expressions such as "According to the passages" or "Based on the passages".

```
{% for passage in passages -%}
- {{passage.text|safe}}
{% endfor %}
```

Question: {{question}}

Please answer directly the question above in {{language}}.

Highlighting Relevant Segments

You are an agent that verifies if sentences are supported by a given context. You will be given a context made of several passages referred as "Passage" split into sentences, each with a numeral identifier, and each referred as "Sentence". Your task is to determine if the sentence is supported by some of the passages (using label 1) or is not supported (using label 0). Please follow this process:

(1) Explain why the sentence is supported or not supported, please write your reasoning in between the tags <rationale></rationale>. If a sentence is supported, list the number of the sentences in the passage or several passages that supported it. If a sentence is not supported, explain why is not supported.

(2) Write the final label for the sentence in between the tags <label></label>.

(3) Write the id of the supporting sentences from passages in between the tags <references></references> separated by comma. Remember to use only 0 or 1 for label.

```
### Context
{% for passage in passages -%}
Passage {{loop.index}}:
{% set outer_loop = loop %}
{% for sent in passage.text %}
    {{outer_loop.index}}.{{loop.index}}:
    {{sent|safe}}
{% endfor %}

{% endfor %}
```

```
### Sentence:
{{sentence|safe}}
```

G Prompts for automated evaluation

The four prompts evaluate answer faithfulness to source passages with increasing complexity: Zero-shot (ZS) provides basic supported/not-supported classification, Chain of Thought (COT) adds explicit reasoning steps, Annotation Guidelines (AG) includes detailed evaluation criteria, and AG+COT combines detailed guidelines with reasoning steps. All prompts output their final classification in <answer> tags.

Zero-shot Prompt (ZS)

You are an automatic annotator tasked with determining whether a given answer is grounded in the list of provided evidence passages. Your role is to carefully analyze the relationship between the answer and the evidence, and then classify the answer as either "Supported" or "Not Supported". Provide your answer directly in <answer> tag.

Evidence Passages:

```
{% for passage in context %}

    . {{loop.index}}: {{passage.text}}

{% endfor %}
```

Answer: {{answer segment}}

Now provided your label directly as "Supported" or "Not Supported".

Chain of thought (COT)

You are an automatic annotator tasked with determining whether a given answer is grounded in the list of provided evidence passages. Your role is to carefully analyze the relationship between the answer and the evidence, write your reasoning in between the tags `<rationale></rationale>` then classify the answer as either "Supported" or "Not Supported" and write answer in `<answer> </answer>` tag

Evidence Passages:

```
{% for passage in context %}
```

```
    {{loop.index}}: {{passage.text}}
```

```
{% endfor %}
```

Answer: {{answer segment}}

Now provided your label directly as "Supported" or "Not Supported".

Annotation Guidelines (AG)

Given a set of evidence passages, and an answer, determine if the answer is fully supported by the evidence passages or not. Analyze each sentence of the answer carefully and verify that all information it contains is explicitly stated in or can be directly inferred from the evidence passages.

Output "Not Supported" if ANY of the following are true:

The answer contains any information not explicitly stated in or directly inferable from the passages.

The answer contradicts any information in the passages.

The answer introduces any new information not found in the passages.

The answer misrepresents or inaccurately paraphrases information from the passages.

The answer draws conclusions not logically supported by the given information.

The answer changes the level of certainty, specificity, or nuance from what is expressed in the passages.

The answer does not directly address the specific aspect asked about in the question.

The answer conflates or misrepresents separate pieces of information when summarizing multiple passages.

Output Supported otherwise.

Provide this determination without any additional explanation in `<answer></answer>` tags. Analyze thoroughly but output only the single-word label.

Evidence Passages:

```
{% for passage in context %}
```

```
    {{loop.index}}: {{passage.text}}
```

```
{% endfor %}
```

Answer: {{answer segment}}

Now provided your label directly as "Supported" or "Not Supported".

Annotation guidelines with Chain of thought (AG + COT)

Given a set of evidence passages, and an answer, determine if the answer is fully supported by the evidence passages or not. Analyze each sentence of the answer carefully and verify that all information it contains is explicitly stated in or can be directly inferred from the evidence passages.

Output **"Not Supported"** if ANY of the following are true:

The answer contains any information not explicitly stated in or directly inferable from the passages.

The answer contradicts any information in the passages.

The answer introduces any new information not found in the passages.

The answer misrepresents or inaccurately paraphrases information from the passages.

The answer draws conclusions not logically supported by the given information.

The answer changes the level of certainty, specificity, or nuance from what is expressed in the passages.

The answer does not directly address the specific aspect asked about in the question.

The answer conflates or misrepresents separate pieces of information when summarizing multiple passages.

Output **Supported** otherwise.

write your reasoning in between the tags **<rationale></rationale>** and Provide your final answer in **<answer></answer>** tags.

Evidence Passages:

```
{% for passage in context %}
  {{loop.index}}: {{passage.text}}
{% endfor %}
```

Answer: {{answer segment}}

Now provided your label directly as "Supported" or "Not Supported".

H More Results on Meta-Evaluators

	EN	DE	ES	FR	HI
GPT4o mini	68.37 \pm 2.84	73.58 \pm 2.86	69.84 \pm 2.26	73.74 \pm 2.29	74.10 \pm 2.61
Qwen 2.5 32B	62.45 \pm 3.00	76.79 \pm 2.76	71.09 \pm 1.98	74.38 \pm 2.28	75.41 \pm 2.72
Llama 3.2 90B	62.60 \pm 2.73	63.11 \pm 2.70	63.36 \pm 2.09	63.17 \pm 1.97	75.05 \pm 2.87
Llama 3.2 11B	60.29 \pm 2.84	59.91 \pm 3.13	59.02 \pm 2.51	65.02 \pm 2.46	65.22 \pm 2.75

Table 10: Balanced Accuracy for five languages and four LLMs using AG+COT prompting strategy. Standard errors are calculated using 1000 Bootstrap iterations.

Prompt	EN	DE	ES	FR	HI
ZS	57.22 \pm 2.14	61.54 \pm 1.98	56.94 \pm 1.26	60.00 \pm 1.39	68.96 \pm 2.22
COT	60.71 \pm 2.34	63.83 \pm 2.15	60.61 \pm 1.44	63.85 \pm 1.53	69.68 \pm 2.42
AG	62.72 \pm 2.24	68.87 \pm 2.27	63.61 \pm 1.36	68.32 \pm 1.53	70.20 \pm 2.02
AG + COT	63.42 \pm 2.44	68.35 \pm 2.37	65.83 \pm 1.64	69.08 \pm 1.71	72.44 \pm 2.07

Table 11: Balanced Accuracy for five languages using four different prompting strategies, averaged across four LLMs (GPT-4o mini, Llama 3.2 90B, Llama 3.2 11B, and Qwen 2.5 32B). Standard errors are calculated using 1000 Bootstrap iterations.

Language	Prompt	GPT4o mini	Llama 3.2 90B	Llama 3.2 11B	Qwen 2.5 32B
EN	ZS	59.8	58.0	51.0	60.1
	COT	61.8	59.1	60.0	62.0
	AG	70.0	59.4	53.0	68.5
	AG + COT	68.4	62.6	60.2	62.5
DE	ZS	61.1	58.6	54.8	71.9
	COT	63.3	60.0	59.6	72.7
	AG	72.0	68.6	60.6	74.6
	AG + COT	73.7	63.3	60.0	76.8
ES	ZS	53.8	56.8	53.7	63.7
	COT	56.2	58.2	61.9	66.2
	AG	69.7	59.6	54.3	70.8
	AG + COT	69.9	63.4	59.1	71.1
FR	ZS	59.5	56.3	56.0	68.5
	COT	60.4	58.5	65.1	71.6
	AG	74.5	62.4	60.3	76.2
	AG + COT	73.7	63.2	65.0	74.4
HI	ZS	72.2	64.3	66.1	73.3
	COT	72.1	67.6	65.4	73.8
	AG	73.8	67.1	65.5	74.5
	AG + COT	74.2	75.1	65.2	75.5

Table 12: Meta evaluation results for all combination of LLMs and prompts we have tested.

Fine-Grained Label	#S	GPT-4o mini	Llama 3.2 90B
Logical conclusion	151	80.1	87.4
Direct paraphrase	182	93.4	97.3
Other	2	50.0	50.0
We. Acc Sup.		87.2	92.5
Adds new information	81	67.9	33.3
Nuance shift	30	26.7	16.7
Mis-referencing	20	35.0	20.0
Contradiction	32	65.6	50.0
Opinion as fact	12	66.7	25.0
Other mistake	19	84.2	52.6
Wrong reasoning	10	80.0	40.0
We. Acc NS.		60.3	33.8
BAcc		73.7	63.2

Table 13: Accuracy of GPT-4o mini and Llama 3.2 90B on various fine-grained labels, when evaluating French language using AG + COT prompt

I Fine grained Analysis

To demonstrate the benchmark’s capability for fine-grained analysis, Table 13 breaks down automatic evaluation performance by error type for two models: GPT-4o mini and Llama 3.2 90B, both evaluating French answers using AG + COT prompt. The upper section of the table shows the distribution of errors when the ground truth class is *Supported*, categorized into logical conclusion, direct paraphrase, or other correct categories. The lower section details the types of mistakes made by each model for the *Not supported* category. We observe that both models show a similar pattern for the *Supported* category, with logical conclusions being the most common (69.8% for GPT-4o mini and 76.0% for Llama 3.2 90B), followed by direct paraphrases. For the *Not supported* category, both models struggle most with detecting “adding new information” (32.1% and 40.0% of mistakes, respectively) and “nuance shifts” (27.2% and 18.5%). The ranking of error types is consistent across both models, despite their different overall BAcc (73.7% for GPT-4o mini vs 63.2% for Llama 3.2 90B see Table 12 in Appendix), suggesting similar challenges in evaluation. The benchmark’s fine-grained labeling system provides a detailed view of evaluation challenges across languages. By categorizing different types of faithfulness violations, it reveals which specific errors are harder to detect for each model, offering insights into both the strengths and limitations of automated evaluation methods.