

The *FACTS Grounding* Leaderboard: Benchmarking LLMs' Ability to Ground Responses to Long-Form Input

Alon Jacovi^{*,♣}, Andrew Wang^{*,♣}, Chris Alberti^{*,◇}, Connie Tao^{*,◇}, Jon Lipovetz^{*,♣}, Kate Olszewska^{*,◇}, Lukas Haas^{*,◇}, Michelle Liu^{*,♡}, Nate Keating^{*,♣}, Adam Bloniarz[♣], Carl Saroufim[♡], Corey Fry[♣], Dror Marcus[♣], Doron Kukliansky[♣], Gaurav Singh Tomar[◇], James Swirhun[♡], Jinwei Xing[♡], Lily Wang[♡], Madhu Gurumurthy[♡], Michael Aaron[♣], Moran Ambar[♣], Rachana Fellingner[♣], Rui Wang[♡], Zizhao Zhang[♡], Sasha Goldshtein[♣] and Dipanjan Das^{*,◇}

^{*}Equal Contribution, [◇]Google DeepMind, [♣]Google Research, [♡]Google Cloud, [♣]Kaggle

We introduce *FACTS Grounding*, an online leaderboard and associated benchmark that evaluates language models' ability to generate text that is factually accurate with respect to given context in the user prompt. In our benchmark, each prompt includes a user request and a full document, with a maximum length of 32k tokens, requiring long-form responses. The long-form responses are required to be fully grounded in the provided context document while fulfilling the user request. Models are evaluated using automated judge models in two phases: (1) responses are disqualified if they do not fulfill the user request; (2) they are judged as accurate if the response is fully grounded in the provided document. The automated judge models were comprehensively evaluated against a held-out test-set to pick the best prompt template, and the final factuality score is an aggregate of multiple judge models to mitigate evaluation bias. The *FACTS Grounding* leaderboard will be actively maintained over time, and contains both public and private splits to allow for external participation while guarding the integrity of the leaderboard. It can be found at <https://www.kaggle.com/facts-leaderboard>.

1. Introduction

Factuality is one of the most challenging aspects of Large Language Models (LLMs), referring to a model's ability to generate factually accurate responses in information-seeking scenarios. Commonly, this area of research can be divided into two distinct scenarios: (1) factuality with respect to *given context*, such as a user request and grounding documents, such that the model response is fully grounded in the input (by this, we imply that a model response has the highest degree of faithfulness to given context as defined by Rashkin et al., 2023), and (2) factuality with respect to *external sources and general world knowledge* (Tang et al., 2024, cf. Pan et al., 2023; Rashkin et al., 2023; Zhao et al., 2024b). There may be subtle use cases that fall in the middle, but most existing literature in the factuality area focus on behaviors that map to these two distinct scenarios. Summarization is a narrow but important example of the first scenario—a summary's various claims should be accurate with respect to the source documents that are being summarized (Bishop et al., 2023; Krishna et al., 2023). Generating short-form, factually accurate answers to factoid questions from an LLM's parametric knowledge is an example use case of the second scenario, as recently discussed by Wei et al. (2024a).

Ensuring factual accuracy while generating LLM responses is challenging. The principal challenges in LLM factuality are *modeling* (i.e., architecture, training, and inference) and *measurement* (i.e., evaluation methodology, data and metrics). On the modeling front, LLMs are inherently difficult to optimize for this goal: typically, LLM pretraining is optimized to predict the next token given previous tokens in the text. While this objective may teach models salient world knowledge, it does not directly optimize the model towards the various factuality scenarios, instead encouraging the model to generate generally *plausible* text. Furthermore, critically, recent research suggests that the

training process (regardless of specific models or training data) inherently admits non-factual text generation (Kalai and Vempala, 2024). Subsequent post-training, via supervised finetuning and reinforcement learning, can steer the model towards improved factuality (Huang and Chen, 2024; Lan et al., 2023; Roit et al., 2023). Other mitigation proposals include inference-time methods, such as prompting or model state interpretability (Lee et al., 2024; Su et al., 2024). However, enhancing factuality through these methods can compromise other desirable attributes, such as creativity and novelty, which are also optimized during these stages. This creates a delicate balance, making it difficult to simultaneously achieve all desired outcomes (Roit et al., 2023).

Due to the above, factuality is expected to remain a research challenge for the foreseeable future. In this work, we focus on the challenge of measuring progress in factuality. Measurement is in itself difficult, due to the different model behaviors that pertain to the aforementioned factuality scenarios. Some settings are more difficult than others: in *long-form generation* tasks, the subject of our benchmark, each claim in the models’ responses should be thoroughly inspected for their accuracy. This is in contrast to tasks that require measuring the factuality of short responses to questions against ground truth (e.g., Wei et al., 2024a). To complicate matters further, long-form generation settings are in themselves varied and numerous. For example, some benchmarks focus on long-form generation in the setting of grounding to external sources (e.g., Wei et al., 2024b; Zhao et al., 2024a; Zhu et al., 2024) while others focus on a particular task or domain, such as summarization to news or biomedical articles (e.g., Ramprasad et al., 2024; Vectara, 2024).

Here, we investigate a benchmark that focuses on scenario (1)—particularly, the measurement of response factuality with respect to a provided context document of length up to 32k tokens, given varied user requests. This setting requires the model to synthesize information derivable from the document while directly addressing the request. While encompassing summarization as a key use case, it extends to a broader range of requests, including fact finding, analyzing and comparing information, and so forth, while being fully grounded to the input document. We believe that this benchmark fills a gap in evaluating a wider variety of model behaviors pertaining to factuality, in comparison to benchmarks that focus on narrower use cases, e.g. summarization alone (Vectara, 2024). (Note however that we do not capture scenario (2) in this work.)

Factuality of long-form responses is difficult to thoroughly measure at scale; particularly automatic evaluation methods continue to be a challenge and is an active research area (Bishop et al., 2023; Chang et al., 2023; Gekhman et al., 2023; Honovich et al., 2022; Jacovi et al., 2024; Karpinska et al., 2024; Kim et al., 2024; Li et al., 2023; Min et al., 2023; Ramprasad and Wallace, 2024; Song et al., 2024; Tang et al., 2024, *inter alia*). A particular limitation of existing automatic evaluation methods is that specialized factuality classifiers are limited to short context windows, or ones that cannot perform reasoning that is required to evaluate factuality behaviors (Jacovi et al., 2024). Here, we rigorously evaluate our automatic evaluators on held-out test data to validate their performance on our task, and use multiple aggregations to mitigate evaluator bias.

We present the *FACTS Grounding* leaderboard and an associated benchmark measuring the ability of LLMs to ground long-form responses to document context up to length 32k tokens given a user request and additional instructions. The benchmark contains 860 public examples (“Open” split) and 859 private examples (“Blind” split) of natural, complex LLM prompts written by humans to evaluate long-form grounded response generation (see examples in Table 1; details in Section 2). The leaderboard reports various LLMs’ performance on this benchmark using an automated factuality score incorporating an eligibility filter to avoid “hacking” the leaderboard metric (Section 3). The results are available in Section 4. The leaderboard will be actively maintained and updated to include new models and their variants.

Table 1 | Examples from *FACTS Grounding* (Public).

System Instruction	Context Document Description	Context Tokens	User Request
Answer the question using only the information provided in the context. Do not rely on external knowledge or sources.	The development and deployment of an autonomous robotic system designed to clean skyscraper windows, highlighting its technological advancements, safety implications, and potential impact on the window-washing industry.	~1.1k	My sister and her dog live in NYC. I’ve visited there and have always been fascinated with their tall buildings. Then I thought...someone has to clean those! Then next thing you know, window washing robots popped up on my feed. How do these robots work? Also what does this mean for the people who do those jobs?
Provide a response based solely on the information provided in the prompt. External sources and prior knowledge must not be used.	Legal interpretations and effects of the medical marijuana appropriations rider on federal marijuana prosecutions, focusing on differing circuit court approaches to determining compliance with state medical marijuana laws.	~1.6k	What did the first circuit conclude?
This task requires you to answer questions based solely on the information provided in the prompt. You are not allowed to use any external resources or prior knowledge. Present your answer in headed sections with an explanation for each section. Each explanation should be in bullet points with exactly three bullet points.	Comparison of different economic systems, including free market, command, and mixed economies, highlighting their key characteristics, advantages, and disadvantages.	~0.9k	which famous economists are mentioned?
Provide your response in a professional and formal tone. Use the information given in the document without referring to external sources or requiring additional context. Avoid using technical jargon or acronyms that are not explained within the document.	Compilation of money-saving tips for college students, categorized into recreation and entertainment, food and basic needs, clothing, budgeting/spending plan, transportation, savings, and conserving resources.	~1.6k	What are some tips on saving money?
Answer the question based solely on the information provided in the passage. Do not use any external knowledge or resources.	A study that investigates the correlation between advanced maternal age (40+) and increased risk of obstetric, fetal, and neonatal complications compared to women aged 25-35.	~2.1k	Researchers at Foch Hospital in France published this study of pregnancy outcomes in two groups of patients. Please summarize outcomes across the three kinds of complications that the researchers studied.

2. Data

Underlying the *FACTS Grounding* leaderboard is a carefully curated collection of documents and associated user requests that were written by human raters and then subjected to thorough validation and filtering. The complete methodology is outlined in the following subsections. Table 1 provides concrete examples of data instances in the collection. To ensure the reliability of the benchmark, both public and private leaderboard splits were constructed using a balanced random sampling strategy.

2.1. Annotation

In order to create our evaluation set, third-party human raters were instructed to design prompts requiring the processing of long-form input and the writing of long-form output. These tasks include Q&A, summarization, and document rewriting. Each example within our evaluation set consists of a context, which is a document or set of reviews sourced from the web, paired with a non-trivial user request that can be addressed using the provided context, necessitating a long-form response. Additionally, each example includes a system instruction directing the model to generate its response exclusively from the given context, without incorporating external knowledge.

To ensure the diversity of the evaluation set, prompts were generated across a range of document




Figure 1 | Distributions of context domain and of task requested by the user as a percent of the total set of prompts in the benchmark.

lengths (up to 32k tokens) and various enterprise domains, including finance, technology, retail, medical, and legal. The annotation instructions were carefully designed to avoid prompts requiring creative responses, expert-level domain knowledge, mathematical or logical reasoning, or meta-analysis of the text, such as tone analysis or interpretation of author intent. The specific distributions of enterprise domains and of tasks requested by users are shown in Figure 1.

2.2. Data Quality Assurance

We manually verified all examples after annotation to discard those that did not align with our instructions. In particular, we ensured that all task instructions required the model to exclusively rely on the provided context, and we further removed creative-writing tasks. We also verified that the user requests were non-trivial and did not require domain expertise, mathematical knowledge, or complex reasoning. We additionally removed documents originating from PDFs where optical character recognition (OCR) rendered them unreadable. After our data quality assurance, the final dataset contained context documents with a mean length of 2.5k tokens and a maximum length of 32k tokens.

On data contamination. As the context documents were collected from the internet, it is possible that they were included in models’ pre-training corpora. Although this is noteworthy, we make the following arguments towards the value of this benchmark:

1. The user requests and system instructions, which instruct specifically to only follow information in the context document, are non-contaminated. Responding to novel requests about non-novel documents is an important use-case of language models, and grounding is integral to it. This is unlike many currently available factuality benchmarks which repurpose academic tasks that have likely been contaminated (Sainz et al., 2024).
2. Our factuality score evaluates a distinct dimension of model performance that is not optimized during pre-training. Specifically, it measures the model’s ability to generate responses grounded exclusively in the provided context. This means that the model must not incorporate external knowledge, even if conflicting with the context document, and should also avoid drawing upon any pre-training knowledge to fulfill the user’s request.
3. As all frontier language models were trained on large corpora of web data, parity is maintained for the purpose of a leaderboard.

Table 2 | Evaluation of different judge models and evaluation prompts on a private test-set (N=402, class ratio 87:13; see §3.1). Chosen prompt template for each model via Macro-F₁ in **bold**.

Judge Model	Prompt Template	Macro-F ₁	Acc.	FPR	FNR	F ₁ (+)	F ₁ (-)
Claude 3.5 Sonnet	Span-level	68.85	77.83	20.97	22.38	85.58	52.13
	Implicit span-level	70.24	83.50	45.16	11.34	90.10	50.37
	Response-level	61.88	83.25	72.58	6.69	90.42	33.33
	JSON	56.04	64.78	33.87	35.47	75.64	36.44
	JSON (alt)	55.37	66.75	46.77	30.81	77.91	32.84
	JSON w. double-check	49.50	54.68	25.81	48.84	65.67	33.33
	SimpleQA template	55.39	85.22	88.71	1.45	91.87	18.92
Gemini 1.5 Pro	Span-level	55.84	79.31	79.03	10.17	88.03	23.64
	Implicit span-level	56.66	85.47	87.10	1.45	91.99	21.33
	Response-level	48.82	82.02	95.16	4.07	90.04	7.59
	JSON	71.47	86.95	56.45	5.23	92.48	50.47
	JSON (alt)	66.03	85.96	69.35	4.07	92.05	40.00
	JSON w. double-check	64.89	76.35	37.10	21.22	84.95	44.83
	SimpleQA template	51.54	84.73	93.55	1.16	91.64	11.43
GPT-4o	Span-level	63.08	81.53	64.52	10.17	89.18	36.97
	Implicit span-level	55.43	83.99	87.10	3.20	91.11	19.75
	Response-level	51.54	84.73	93.55	1.16	91.64	11.43
	JSON	69.68	80.54	32.26	17.15	87.83	51.53
	JSON (alt)	66.78	82.02	53.23	11.63	89.28	44.27
	JSON w. double-check	57.62	64.04	17.74	39.24	74.11	41.13
	SimpleQA template	47.04	83.74	98.39	1.45	91.13	2.94

3. Metrics

The final factuality score in *FACTS Grounding* is calculated as an aggregation of factuality verdicts from multiple judge models after the disqualification of ineligible responses that do not sufficiently succeed at responding to the user request.

3.1. Unadjusted Factuality Score

The principal component of our evaluation process is an unadjusted *factuality score*.

First, we utilize a language model judge to produce a binary classification label identifying whether a full model response is grounded in the user request and the context document given an instruction (see Table 1). A model response is marked with a positive label (“accurate”) if all the claims in the response are grounded in the contents of the prompt, or do not require grounding; the response is marked with a negative label (“not accurate”) if a single claim that bears information is deemed to be not grounded in the contents of the prompt. We use three different judge models in order to reduce the bias of a particular judge model, as models have been shown to be biased towards favorably judging their own outputs (Wataoka et al., 2024). The judges are prompted language models in all cases—*Gemini 1.5 Pro* (Gemini Team: R. Anil et al., 2023), *GPT-4o* (Achiam et al., 2023), and *Claude 3.5 Sonnet* (Anthropic, 2024). To select the right judge prompt, we tested seven different prompt templates and evaluated them based on alignment with human judgements on a held-out private set of prompts and model responses (N=402) which were annotated with golden labels. The prompt templates used in the evaluation are in Appendix A, and results are shown in Table 2. For each judge model, we select the best-performing prompt template via Macro-F₁.

Given the three judges, the individual factuality score for each judge is the percentage of accurate responses, and the unadjusted factuality score is the average of all judge models’ scores.

Table 3 | Examples of ineligible responses: these responses, while fully grounded in the context document, fail to address the user request meaningfully and are consequently considered ineligible.

Context Document Description	User Request	Ineligible Response	Rationale
A research report on renewable energy sources, including wind, solar, and hydroelectric power, with specific statistics and case studies.	Can you summarize the key advantages and disadvantages of wind energy from this document?	Wind energy is good because it’s renewable and clean, but it has some challenges too.	(1) The response is extremely vague, failing to provide any detailed or specific points from the document, such as the cost-effectiveness, geographic limitations, or impacts on wildlife. (2) It doesn’t engage with the query’s focus on "key" advantages and disadvantages.
A company’s annual financial report, discussing quarterly earnings, expenditures, future investments, and an analysis of the market environment.	Summarize the main reasons the company’s revenue decreased in Q3.	The company faced challenges in Q3 that impacted its revenue.	(1) The response avoids specifying any reasons, such as market trends, increased competition, or operational setbacks, which would likely be in the document. (2) It doesn’t demonstrate an attempt to engage with or extract relevant details.
A historical article on the causes and consequences of the Great Depression.	What were the main causes of the Great Depression as explained in the document?	The Great Depression was a difficult time in history with many causes and effects.	(1) The response provides no substantive information on the causes, such as stock market speculation, bank failures, or trade policies, which were discussed in the document. (2) It ignores the user’s explicit focus on the "main causes."

3.2. Disqualifying Ineligible Responses

Metrics that are focused on evaluating the factuality of the generated text, with respect to a context document or otherwise, can be circumvented by ignoring the intent behind the user request. By giving shorter responses that evade conveying comprehensive information, even if such content was an important aspect of a user request, it is possible to achieve a high factuality score while not providing a helpful response. See illustrative examples in Table 3.

We safeguard against such responses by detecting them with prompted LLM judges that use the same base models as in Section 3.1, with prompt templates described in Appendix A. At a response level, similar to unadjusted factuality score, we treat instruction-following as a distinct task for an LLM judge. This task involves a binary classification of each model response, determining whether it sufficiently addresses the user’s request. Each response is assigned a binary label indicating its eligibility: either “eligible,” signifying that it answers the user’s request, or “ineligible,” otherwise. Ineligible responses are disqualified from factuality evaluation and the final factuality score is adjusted such that ineligible responses are deemed as *inaccurate*.

We investigate prompting each LLM judge with a prompt that contains either the user request only or the user request and the context document. We evaluate these two prompt templates on a separate test set which includes golden labels for instruction-following eligibility and assess whether the predicted ratings align with these golden labels (see Table 4 for results). For each LLM judge, we select the prompt with the highest Macro-F₁. Per response, eligibility classifications across the three judges are ensembled by consensus. A response is considered ineligible only if all three judges consider the response ineligible. Ensembling via consensus focuses this benchmark on evaluating grounding while still filtering out the worst quality responses.

Table 4 | Evaluation of prompt template for the eligible responses detection task on a private test-set (N=450; see §3.2). Chosen prompt template for each model via Macro-F₁ in **bold**.

Judge Model	Prompt Template	Macro-F ₁	Acc.	FPR	FNR	F ₁ (+)	F ₁ (-)
Claude 3.5 Sonnet	User request only	60.88	68.22	16.33	62.67	43.92	77.83
	User request and context document	58.14	69.56	8.67	74.00	36.28	80.00
Gemini 1.5 Pro	User request only	56.28	67.11	12.33	74.00	34.51	78.04
	User request and context document	47.56	68.44	1.33	92.00	14.46	80.65
GPT-4o	User request only	55.16	69.56	5.33	80.67	29.74	80.57
	User request and context document	50.55	69.56	1.33	88.67	19.88	81.21

Table 5 | *FACTS Grounding* results representing unadjusted *factuality score* (no disqualification). Results are reported with a 95% confidence interval.

Response Model	Judge Models						Average	Fused Rank
	Open (N=860)			Blind (N=859)				
	Gemini 1.5 Pro	GPT-4o	Claude 3.5 Sonnet	Gemini 1.5 Pro	GPT-4o	Claude 3.5 Sonnet		
Gemini 1.5 Flash	91.4(±1.9)	82.0(±2.6)	84.9(±2.4)	90.7(±1.9)	80.7(±2.6)	85.1(±2.4)	85.8 (±1.7)	1
Gemini 2.0 Flash Experimental	88.7(±2.1)	79.5(±2.7)	86.7(±2.3)	91.5(±1.9)	79.3(±2.7)	88.0(±2.2)	85.6(±1.7)	2
Gemini 1.5 Pro	90.2(±2.0)	76.0(±2.9)	83.6(±2.5)	89.8(±2.0)	73.2(±3.0)	83.5(±2.5)	82.7(±1.8)	3
Claude 3.5 Sonnet	88.4(±2.1)	72.7(±3.0)	87.7(±2.2)	88.4(±2.1)	73.8(±2.9)	82.2(±2.6)	82.2(±1.8)	4
GPT-4o	86.7(±2.3)	71.7(±3.0)	79.2(±2.7)	88.0(±2.2)	75.9(±2.9)	77.4(±2.8)	79.8(±1.9)	5
Claude 3.5 Haiku	85.8(±2.3)	67.2(±3.1)	82.0(±2.6)	80.1(±2.7)	62.5(±3.2)	74.3(±2.9)	75.3(±2.0)	6
GPT-4o mini	80.8(±2.6)	62.1(±3.2)	71.4(±3.0)	83.0(±2.5)	65.0(±3.2)	70.7(±3.0)	72.2(±2.1)	7
OpenAI o1-mini	70.8(±3.0)	50.2(±3.3)	63.1(±3.2)	75.1(±2.9)	51.3(±3.3)	64.6(±3.2)	62.5(±2.3)	8
OpenAI o1-preview	69.2(±3.1)	50.1(±3.3)	65.3(±3.2)	70.0(±3.1)	53.2(±3.3)	65.0(±3.2)	62.1(±2.3)	9

4. Results

Table 5 and Table 6 contain the unadjusted and final *factuality scores* before and after disqualifying ineligible responses, respectively. The tested models are *Gemini 1.5 Pro* and *Flash* (Gemini Team: R. Anil et al., 2023), *Gemini 2.0 Flash Experimental* (Gemini Team, 2024), *GPT-4o* (Achiam et al., 2023), *OpenAI o1-preview* and *o1-mini* (OpenAI, 2024), *Claude 3.5 Haiku* and *Sonnet* (Anthropic, 2024). For the "Fused Rank" metric, we employ a ranking aggregation method that combines the six individual model rankings—derived from each split and judge model—into a single, unified ranking using the Condorcet algorithm. The resulting fused rank exactly aligns with the ranking obtained using the final factuality score.

On aggregating multiple judge models. Consistent with prior research (Liu et al., 2024; Wataoka et al., 2024; Xu et al., 2024; Ye et al., 2024; Zheng et al., 2023), we found that models generally rate their own outputs higher than those of other models, exhibiting a mean increase of +3.23%. While the use of multiple judge models increases the computational cost of evaluation, this approach is essential when the judge models themselves are also subject to evaluation, as is the case in our work.

On ineligible responses. Disqualifying ineligible responses leads to a reduction of 1%–5% in the final factuality score, as these responses are treated as inaccurate. The disqualification process also induces a minor shift in model rankings; for instance, *Gemini 1.5 Flash* moves from rank 1 to rank 2.

Table 6 | *FACTS Grounding* Results representing the final *factuality score* after disqualifying ineligible responses (see §3.2). Results are reported with a 95% confidence interval.

Response Model	Judge Models						Average	Fused Rank
	Open (N=860)			Blind (N=859)				
	<i>Gemini 1.5 Pro</i>	<i>GPT-4o</i>	<i>Claude 3.5 Sonnet</i>	<i>Gemini 1.5 Pro</i>	<i>GPT-4o</i>	<i>Claude 3.5 Sonnet</i>		
Gemini 2.0 Flash Experimental	86.4(±2.3)	77.4(±2.8)	84.8(±2.4)	89.3(±2.1)	77.2(±2.8)	86.3(±2.3)	83.6(±1.8)	1
Gemini 1.5 Flash	88.1(±2.2)	79.2(±2.7)	82.6(±2.5)	87.3(±2.2)	77.9(±2.8)	82.3(±2.6)	82.9(±1.8)	2
Gemini 1.5 Pro	87.4(±2.2)	73.8(±2.9)	81.4(±2.6)	86.5(±2.3)	70.2(±3.1)	80.8(±2.6)	80.0(±1.9)	3
Claude 3.5 Sonnet	87.3(±2.2)	71.7(±3.0)	86.7(±2.3)	82.0(±2.6)	67.8(±3.1)	81.0(±2.6)	79.4(±1.9)	4
GPT-4o	85.3(±2.4)	70.7(±3.0)	78.5(±2.7)	86.5(±2.3)	74.9(±2.9)	76.8(±2.8)	78.8(±1.9)	5
Claude 3.5 Haiku	84.8(±2.4)	66.6(±3.2)	81.2(±2.6)	77.9(±2.8)	61.0(±3.3)	73.7(±2.9)	74.2(±2.1)	6
GPT-4o mini	79.4(±2.7)	60.8(±3.3)	70.7(±3.0)	81.4(±2.6)	63.7(±3.2)	70.1(±3.1)	71.0(±2.1)	7
OpenAI o1-mini	70.2(±3.1)	49.8(±3.3)	62.7(±3.2)	74.2(±2.9)	50.8(±3.3)	64.3(±3.2)	62.0(±2.3)	8
OpenAI o1-preview	68.6(±3.1)	49.7(±3.3)	65.0(±3.2)	69.4(±3.1)	52.6(±3.3)	64.6(±3.2)	61.7(±2.3)	9

5. Conclusion

The *FACTS Grounding* leaderboard is designed to rigorously challenge language models’ ability to maintain factual accuracy when generating long-form responses grounded in a document provided within the prompt, and in accordance with a user’s specific request and instructions. We encourage other researchers to utilize this benchmark to advance both the factual capabilities of models and the methodologies for evaluating factuality.

6. Contributions and Acknowledgements

- **Experimental design:** Alon Jacovi, Andrew Wang, Chris Alberti, Jon Lipovetz, and Michelle Liu created the experimental design behind the benchmark and ran all the reported experiments.
- **Organization:** Connie Tao, Dipanjan Das, Kate Olszewska, Lukas Haas, and Nate Keating managed the overall organization of the effort from start to completion.
- **Early experimentation:** Adam Bloniarz, Carl Saroufim, Corey Fry, Dror Marcus, Doron Kukliansky, Gaurav Singh Tomar, James Swirhun, Jinwei Xing, Lily Wang, Madhu Gurumurthy, Michael Aaron, Moran Ambar, Rachana Fellingner, Rui Wang, Zizhao Zhang and Sasha Goldshtein contributed to ideas, conducted data collection and ran early experiments.
- **Sponsors:** Avinatan Hassidim, D. Sculley, Fernando Pereira, Koray Kavukcuoglu, Slav Petrov, Ya Xu, and Yossi Matias sponsored the effort and provided technical guidance.
- Slav Petrov and Madhu Gurumurthy proposed the idea for this leaderboard.

All authors wrote parts of the report.

We would also like to thank:

- **Gemini team** for the support and model access.
- **Kaggle team** for their expertise and releasing the leaderboard.
- **Expert data annotators** who helped to collect examples in the paper.
- Our reviewers **Kellie Webster** and **Phoebe Kirk** for valuable feedback.

References

- J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Anthropic. The Claude 3 model family: Opus, Sonnet, Haiku, 2024. URL https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf.
- J. A. Bishop, Q. Xie, and S. Ananiadou. LongDocFACTScore: Evaluating the factuality of long document abstractive summarisation. *arXiv preprint arXiv:2309.12455*, 2023.
- Y. Chang, K. Lo, T. Goyal, and M. Iyyer. Boookscore: A systematic exploration of book-length summarization in the era of LLMs. *arXiv preprint arXiv:2310.00785*, 2023.
- Z. Gekhman, J. Herzig, R. Aharoni, C. Elkind, and I. Szpektor. Trueteacher: Learning factual consistency evaluation with large language models. *arXiv preprint arXiv:2305.11171*, 2023.
- Gemini Team. Introducing gemini 2.0: our new ai model for the agentic era, 2024. URL <https://blog.google/technology/google-deepmind/google-gemini-ai-update-december-2024>. Accessed: 2024-12-11.
- Gemini Team: R. Anil, S. Borgeaud, J.-B. Alayrac, J. Yu, R. Soricut, J. Schalkwyk, A. M. Dai, A. Hauth, K. Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- O. Honovich, R. Aharoni, J. Herzig, H. Taitelbaum, D. Kukliansy, V. Cohen, T. Scialom, I. Szpektor, A. Hassidim, and Y. Matias. TRUE: Re-evaluating factual consistency evaluation. *arXiv preprint arXiv:2204.04991*, 2022.
- C.-W. Huang and Y.-N. Chen. FactAlign: Long-form factuality alignment of large language models. *arXiv preprint arXiv:2410.01691*, 2024.
- A. Jacovi, M. Ambar, E. Ben-David, U. Shaham, A. Feder, M. Geva, D. Marcus, and A. Caciularu. Coverbench: A challenging benchmark for complex claim verification. *arXiv preprint arXiv:2408.03325*, 2024.
- A. T. Kalai and S. S. Vempala. Calibrated language models must hallucinate. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 160–171, 2024.
- M. Karpinska, K. Thai, K. Lo, T. Goyal, and M. Iyyer. One thousand and one pairs: A "novel" challenge for long-context language models. *arXiv preprint arXiv:2406.16264*, 2024.
- Y. Kim, Y. Chang, M. Karpinska, A. Garimella, V. Manjunatha, K. Lo, T. Goyal, and M. Iyyer. FABLES: Evaluating faithfulness and content selection in book-length summarization. *arXiv preprint arXiv:2404.01261*, 2024.
- K. Krishna, E. Bransom, B. Kuehl, M. Iyyer, P. Dasigi, A. Cohan, and K. Lo. LongEval: Guidelines for human evaluation of faithfulness in long-form summarization. *arXiv preprint arXiv:2301.13298*, 2023.
- Z. Lan, W. Li, J. Su, X. Xiao, J. Liu, W. Wu, and Y. Lyu. Factgen: Faithful text generation by factuality-aware pre-training and contrastive ranking fine-tuning. *Journal of Artificial Intelligence Research*, 76:1281–1303, 2023.

- S. Lee, H. Hsu, and C.-F. Chen. LLM hallucination reasoning with zero-shot knowledge test. *arXiv preprint arXiv:2411.09689*, 2024.
- J. Li, X. Cheng, W. X. Zhao, J.-Y. Nie, and J.-R. Wen. HaluEval: A large-scale hallucination evaluation benchmark for large language models. *arXiv preprint arXiv:2305.11747*, 2023.
- Y. Liu, N. Moosavi, and C. Lin. LLMs as narcissistic evaluators: When ego inflates evaluation scores. In L.-W. Ku, A. Martins, and V. Srikumar, editors, *Findings of the Association for Computational Linguistics: ACL 2024*, pages 12688–12701, Bangkok, Thailand, Aug. 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.753. URL <https://aclanthology.org/2024.findings-acl.753>.
- S. Min, K. Krishna, X. Lyu, M. Lewis, W.-t. Yih, P. W. Koh, M. Iyyer, L. Zettlemoyer, and H. Hajishirzi. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation. *arXiv preprint arXiv:2305.14251*, 2023.
- OpenAI. Learning to reason with LLMs, 2024. URL <https://openai.com/index/learning-to-reason-with-llms>.
- L. Pan, X. Wu, X. Lu, A. T. Luu, W. Y. Wang, M.-Y. Kan, and P. Nakov. Fact-checking complex claims with program-guided reasoning. *arXiv preprint arXiv:2305.12744*, 2023.
- S. Ramprasad and B. C. Wallace. Do automatic factuality metrics measure factuality? A critical evaluation. *arXiv preprint arXiv:2411.16638*, 2024.
- S. Ramprasad, K. Krishna, Z. C. Lipton, and B. C. Wallace. Evaluating the factuality of zero-shot summarizers across varied domains. *arXiv preprint arXiv:2402.03509*, 2024.
- H. Rashkin, V. Nikolaev, M. Lamm, L. Aroyo, M. Collins, D. Das, S. Petrov, G. S. Tomar, I. Turc, and D. Reitter. Measuring attribution in natural language generation models. *Computational Linguistics*, 49(4):777–840, 2023.
- P. Roit, J. Ferret, L. Shani, R. Aharoni, G. Cideron, R. Dadashi, M. Geist, S. Girgin, L. Hussenot, O. Keller, et al. Factually consistent summarization via reinforcement learning with textual entailment feedback. *arXiv preprint arXiv:2306.00186*, 2023.
- O. Sainz, I. García-Ferrero, A. Jacovi, J. A. Campos, Y. Elazar, E. Agirre, Y. Goldberg, W.-L. Chen, J. Chim, L. Choshen, et al. Data contamination report from the 2024 CONDA shared task. *arXiv preprint arXiv:2407.21530*, 2024.
- Y. Song, Y. Kim, and M. Iyyer. VERIScore: Evaluating the factuality of verifiable claims in long-form text generation. *arXiv preprint arXiv:2406.19276*, 2024.
- W. Su, C. Wang, Q. Ai, Y. Hu, Z. Wu, Y. Zhou, and Y. Liu. Unsupervised real-time hallucination detection based on the internal states of large language models. *arXiv preprint arXiv:2403.06448*, 2024.
- L. Tang, P. Laban, and G. Durrett. MiniCheck: Efficient fact-checking of LLMs on grounding documents. *arXiv preprint arXiv:2404.10774*, 2024.
- Vectara. Hallucination evaluation model (revision 7437011), 2024. URL https://huggingface.co/vectara/hallucination_evaluation_model.
- K. Wataoka, T. Takahashi, and R. Ri. Self-preference bias in LLM-as-a-judge. *arXiv preprint arXiv:2410.21819*, 2024.

- J. Wei, N. Karina, H. W. Chung, Y. J. Jiao, S. Papay, A. Glaese, J. Schulman, and W. Fedus. Measuring short-form factuality in large language models. *arXiv preprint arXiv:2411.04368*, 2024a.
- J. Wei, C. Yang, X. Song, Y. Lu, N. Hu, D. Tran, D. Peng, R. Liu, D. Huang, C. Du, et al. Long-form factuality in large language models. *arXiv preprint arXiv:2403.18802*, 2024b.
- W. Xu, G. Zhu, X. Zhao, L. Pan, L. Li, and W. Wang. Pride and prejudice: LLM amplifies self-bias in self-refinement. In L.-W. Ku, A. Martins, and V. Srikumar, editors, *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15474–15492, Bangkok, Thailand, Aug. 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.acl-long.826. URL <https://aclanthology.org/2024.acl-long.826>.
- J. Ye, Y. Wang, Y. Huang, D. Chen, Q. Zhang, N. Moniz, T. Gao, W. Geyer, C. Huang, P.-Y. Chen, N. V. Chawla, and X. Zhang. Justice or prejudice? quantifying biases in llm-as-a-judge. *arXiv preprint arXiv:2410.02736*, 2024.
- W. Zhao, T. Goyal, Y. Y. Chiu, L. Jiang, B. Newman, A. Ravichander, K. Chandu, R. L. Bras, C. Cardie, Y. Deng, et al. WildHallucinations: Evaluating long-form factuality in LLMs with real-world entity queries. *arXiv preprint arXiv:2407.17468*, 2024a.
- Y. Zhao, J. Zhang, I. Chern, S. Gao, P. Liu, J. He, et al. FELM: Benchmarking factuality evaluation of large language models. *Advances in Neural Information Processing Systems*, 36, 2024b.
- L. Zheng, W.-L. Chiang, Y. Sheng, S. Zhuang, Z. Wu, Y. Zhuang, Z. Lin, Z. Li, D. Li, E. Xing, H. Zhang, J. E. Gonzalez, and I. Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 46595–46623. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/91f18a1287b398d378ef22505bf41832-Paper-Datasets_and_Benchmarks.pdf.
- Z. Zhu, Y. Yang, and Z. Sun. HaluEval-Wild: Evaluating hallucinations of language models in the wild. *arXiv preprint arXiv:2403.04307*, 2024.

A. Judge Prompt Templates

Factuality Score.

Response-level

Your task is to check if the Response is accurate to the Evidence. Generate 'Accurate' if the Response is accurate when verified according to the Evidence, or 'Inaccurate' if the Response is inaccurate (contradicts the evidence) or cannot be verified.

```

**Query**:\n\n{user_query}\n\n**End of Query**\n
**Evidence**:\n\n{context}\n\n**End of Evidence**\n
**Response**:\n\n{response}\n\n**End of Response**\n
Let's think step-by-step.
    
```

JSON (Alt)

You are a helpful and harmless AI assistant. You will be provided with a textual context and a model-generated response. Your task is to analyze the response sentence by sentence and classify each sentence according to its relationship with the provided context.

****Instructions:****

- **Decompose the response into individual sentences.****
- **For each sentence, assign one of the following labels:****
 - * ****`supported`****: The sentence is entailed by the given context. Provide a supporting excerpt from the context.
 - * ****`unsupported`****: The sentence is not entailed by the given context. Provide an excerpt that is close but does not fully support the sentence.
 - * ****`contradictory`****: The sentence is falsified by the given context. Provide a contradicting excerpt from the context.
 - * ****`no_rad`****: The sentence does not require factual attribution (e.g., opinions, greetings, questions, disclaimers). No excerpt is needed for this label.
- **For each label, provide a short rationale explaining your decision.**** The rationale should be separate from the excerpt.

****Input Format:****

The input will consist of two parts, clearly separated:

- * ****Context:**** The textual context used to generate the response.
- * ****Response:**** The model-generated response to be analyzed.

****Output Format:****

For each sentence in the response, output a JSON object with the following fields:

- * ****`sentence`****: The sentence being analyzed.
- * ****`label`****: One of ``supported``, ``unsupported``, ``contradictory``, or ``no_rad``.
- * ****`rationale`****: A brief explanation for the assigned label.
- * ****`excerpt`****: A relevant excerpt from the context. Only required for ``supported``, ``unsupported``, and ``contradictory`` labels.

Output each JSON object on a new line.

****Example:****

****Input:****

```

...
Context: Apples are red fruits. Bananas are yellow fruits.

Response: Apples are red. Bananas are green. Enjoy your fruit!
...

**Output:**

{"sentence": "Apples are red.", "label": "supported", "rationale": "The context explicitly states that apples are red.", "excerpt": "Apples are red fruits."}
{"sentence": "Bananas are green.", "label": "contradictory", "rationale": "The context states that bananas are yellow, not green.", "excerpt": "Bananas are yellow fruits."}
{"sentence": "Enjoy your fruit!", "label": "no_rad", "rationale": "This is a general expression and does not require factual attribution.", "excerpt": null}

**Now, please analyze the following context and response:**

**User Query:**
{user_query}

**Context:**
{context}

**Response:**
{response}

```

JSON

```

You are a helpful and harmless AI assistant. You will be provided with a textual context and a model-generated response. Your task is to analyze the response sentence by sentence and classify each sentence according to its relationship with the provided context.

**Instructions:**

1. Decompose the response into individual sentences.
2. For each sentence, assign one of the following labels:
   * `supported`: The sentence is entailed by the given context. Provide a supporting excerpt from the context. The supporting excerpt must fully entail the sentence. If you need to cite multiple supporting excerpts, simply concatenate them.
   * `unsupported`: The sentence is not entailed by the given context. No excerpt is needed for this label.
   * `contradictory`: The sentence is falsified by the given context. Provide a contradicting excerpt from the context.
   * `no_rad`: The sentence does not require factual attribution (e.g., opinions, greetings, questions, disclaimers). No excerpt is needed for this label.
3. For each label, provide a short rationale explaining your decision. The rationale should be separate from the excerpt.
4. Be very strict with your `supported` and `contradictory` decisions. Unless you can find straightforward, indisputable evidence excerpts in the context that a sentence is supported or contradictory, consider it unsupported. You should not employ world knowledge unless it is truly trivial.

**Input Format:**

The input will consist of two parts, clearly separated:

* Context: The textual context used to generate the response.
* Response: The model-generated response to be analyzed.

```

```

**Output Format:**

For each sentence in the response, output a JSON object with the following fields:

* `sentence`: The sentence being analyzed.
* `label`: One of `supported`, `unsupported`, `contradictory`, or `no_rad`.
* `rationale`: A brief explanation for the assigned label.
* `excerpt`: A relevant excerpt from the context. Only required for `supported` and `contradictory` labels.

Output each JSON object on a new line.

**Example:**

**Input:**
...
Context: Apples are red fruits. Bananas are yellow fruits.

Response: Apples are red. Bananas are green. Bananas are cheaper than apples. Enjoy your fruit!
...

**Output:**

{"sentence": "Apples are red.", "label": "supported", "rationale": "The context explicitly states that apples are red.", "excerpt": "Apples are red fruits."}
{"sentence": "Bananas are green.", "label": "contradictory", "rationale": "The context states that bananas are yellow, not green.", "excerpt": "Bananas are yellow fruits."}
{"sentence": "Bananas are cheaper than apples.", "label": "unsupported", "rationale": "The context does not mention the price of bananas or apples.", "excerpt": null}
{"sentence": "Enjoy your fruit!", "label": "no_rad", "rationale": "This is a general expression and does not require factual attribution.", "excerpt": null}

**Now, please analyze the following context and response:**

**User Query:**
{user_query}

**Context:**
{context}

**Response:**
{response}

```

JSON w. double-check (this template uses the *JSON* template initially, while using the below template to double-check each sentence-level classification)

```

Your task is to verify whether a given sentence is entailed by a given context or not. Answer only in YES or NO without any additional text. Do not try to avoid answering, or apologize, or give any answer that isn't simply YES or NO.

**Sentence**
{json_dict["sentence"]}

**Context**
{json_dict["excerpt"]}

```

Span-level

Your task is to check if a specific Span is accurate to the Evidence.
 Generate 'Accurate' if the Span is accurate when verified according to the Evidence or when there is nothing to verify in the Span.
 Generate 'Inaccurate' if the Span is inaccurate (contradicts the evidence), or cannot be verified.

```

**Query**:\n\n{user_query}\n\n**End of Query**\n
**Evidence**\n\n{context}\n\n**End of Evidence**\n
**Response**:\n\n{response}\n\n**End of Response**\n
    
```

You are currently verifying **Span {ix+1}** from the Response.
Span {ix+1}:\n\n{span}\n\n**End of Span {ix+1}**

Is Span {ix+1} accurate or inaccurate when verified according to the Evidence? Point to where in the evidence justifies your answer.

Implicit span-level

Your task is to check if the Response is accurate to the Evidence.
 Generate 'Accurate' if the Response is accurate when verified according to the Evidence, or 'Inaccurate' if the Response is inaccurate (contradicts the evidence) or cannot be verified.

```

**Query**:\n\n{user_query}\n\n**End of Query**\n
**Evidence**\n\n{context}\n\n**End of Evidence**\n
**Response**:\n\n{response}\n\n**End of Response**\n
    
```

Break down the Response into sentences and classify each one separately, then give the final answer: If even one of the sentences is inaccurate, then the Response is inaccurate.

For example, your output should be of this format:
 Sentence 1: <Sentence 1>
 Sentence 1 label: Accurate/Inaccurate (choose 1)
 Sentence 2: <Sentence 2>
 Sentence 2 label: Accurate/Inaccurate (choose 1)
 Sentence 3: <Sentence 3>
 Sentence 3 label: Accurate/Inaccurate (choose 1)
 [...]
 Final Answer: Accurate/Inaccurate (choose 1)

Ineligible responses filter.

The following template is used for both the *user-request-only* and the *user-request-with-context-document* methods. The only difference is whether only the user request or the full prompt are inserted respectively in place of the `user_request_or_full_prompt` placeholder. Responses are ineligible if they contain "Major Issue(s)" in instruction following according to all LLM judges.

```

# Rubrics
Your mission is to judge the response from an AI model, the *test* response, calibrating your judgement using a *baseline* response.
Please use the following rubric criteria to judge the responses:

<START OF RUBRICS>
Your task is to analyze the test response based on the criterion of "Instruction Following". Start your analysis with "Analysis".

**Instruction Following**
Please first list the instructions in the user query.
    
```

In general, an instruction is VERY important if it is specifically asked for in the prompt and deviates from the norm. Please highlight such specific keywords. You should also derive the task type from the user query and include the task-specific implied instructions.

Sometimes, no instruction is available in the user query. It is your job to infer if the instruction is to autocomplete the user query or is asking the LLM for follow-ups.

After listing the instructions, you should rank them in order of importance. After that, INDEPENDENTLY check if the test response and the baseline response meet each of the instructions.

You should itemize, for each instruction, whether the response meets, partially meets, or does not meet the requirement, using reasoning. You should start reasoning first before reaching a conclusion about whether the response satisfies the requirement. Citing examples while reasoning is preferred.

Reflect on your answer and consider the possibility that you are wrong. If you are wrong, explain clearly what needs to be clarified, improved, or changed in the rubric criteria and guidelines.

In the end, express your final verdict as one of the following three json objects:

```

```json
{{
 "Instruction Following": "No Issues"
}}
...

```json
{{
  "Instruction Following": "Minor Issue(s)"
}}
...

```json
{{
 "Instruction Following": "Major Issue(s)"
}}
...

```

<END OF RUBRICS>

```

Your task
User query
<|begin_of_query|>
{user_request_or_full_prompt}
<|end_of_query|>

Test Response:
<|begin_of_test_response|>
{test_response}
<|end_of_test_response|>

Baseline Response:
<|begin_of_baseline_response|>
{baseline_response}
<|end_of_baseline_response|>

```

Please write your analysis and final verdict for the test response.