

✿ LUQ: Long-text Uncertainty Quantification for LLMs

Caiqi Zhang¹, Fangyu Liu^{1*}, Marco Basaldella^{2†}, Nigel Collier¹

¹Language Technology Lab, University of Cambridge ²Amazon Alexa

{cz391, f1399, nhc30}@cam.ac.uk

mbbasald@amazon.co.uk

Abstract

Large Language Models (LLMs) have demonstrated remarkable capability in a variety of NLP tasks. However, LLMs are also prone to generate nonfactual content. Uncertainty Quantification (UQ) is pivotal in enhancing our understanding of a model’s confidence on its generation, thereby aiding in the mitigation of nonfactual outputs. Existing research on UQ predominantly targets short text generation, typically yielding brief, word-limited responses. However, real-world applications frequently necessitate much longer responses. Our study first highlights the limitations of current UQ methods in handling long text generation. We then introduce LUQ with its two variations: LUQ-ATOMIC and LUQ-PAIR, a series of novel sampling-based UQ approaches specifically designed for long text. Our findings reveal that LUQ outperforms existing baseline methods in correlating with the model’s factuality scores (negative coefficient of -0.85 observed for Gemini Pro). To further improve the factuality of LLM responses, we propose LUQ-ENSEMBLE, a method that ensembles responses from multiple models and selects the response with the lowest uncertainty. The ensembling method greatly improves the response factuality upon the best standalone LLM.¹

1 Introduction

Large Language Models (LLMs) have demonstrated significant prowess across a wide range of NLP tasks and are increasingly being used in various downstream applications (Zhao et al., 2023; Chang et al., 2023). However, existing LLMs are susceptible to hallucination, often resulting in the generation of nonfactual or fabricated content (Manakul et al., 2023; Zhang et al., 2023). One way to predict the factuality of an LLM’s output without resorting to resource-intensive fact-checking

procedures is by examining its uncertainty over a user query. Moreover, accurate measurement of a model’s confidence in its generated responses can enable the rejection of answers with high uncertainty, potentially reducing hallucinations and improving the factuality of the output (Geng et al., 2023; Wang et al., 2023).

Although Uncertainty Quantification (UQ) is a well-researched area in machine learning (Gawlikowski et al., 2023), its application in the context of LLMs remains under-explored. One primary limitation is that previous studies on UQ mostly require access to a model’s internal states (e.g., logits) (Murray and Chiang, 2018; Kuhn et al., 2023; Vazhentsev et al., 2023; Duan et al., 2023). However, many best-performing LLMs, such as GPT-4 (OpenAI, 2023), Gemini 1.0 Pro (Gemini Team, 2023), and Claude 2.1 (Anthropic, 2023), are closed-source and only accessible via API calls. This limits the ability to directly analyze their internal processes. Another challenge is that existing research on modeling uncertainty predominantly focuses on short responses, typically less than 10 words in length (Kuhn et al., 2023; Duan et al., 2023; Lin et al., 2023). This is in stark contrast to the more common use cases of LLMs, where responses to queries often far exceed this length, sometimes reaching hundreds of words. Such disparity points to a need for new UQ methods tailored for long-form text generated by LLMs. Therefore, in this study we aim to answer the following research questions: **RQ1:** Are existing UQ methods still effective in the context of long-text generation? **RQ2:** If not, how can we effectively quantify LLMs’ uncertainty for long-form answers? **RQ3:** In what ways can uncertainty scores be utilized to enhance the factuality of model outputs?

We explore UQ for long-text generation (at least 100 words), with an emphasis on using factuality as the key metric of the models’ performance. The main contributions of this paper are:

*Now at Google DeepMind.

†Work done outside of Amazon.

¹<https://github.com/caiqizh/LUQ>

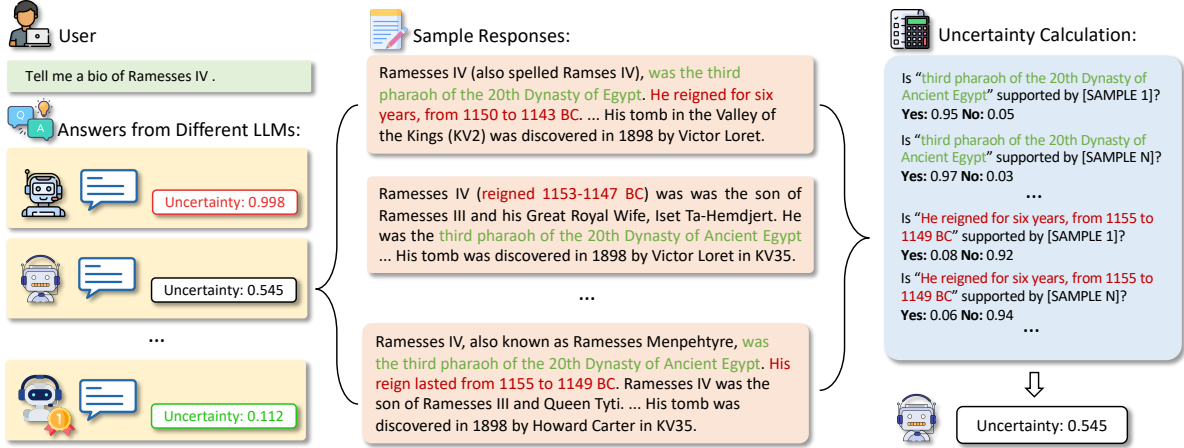


Figure 1: The illustration of the LUQ and LUQ-ENSEMBLE framework. Given a question, various LLMs exhibit differing levels of uncertainty. We generate n sample responses from each LLM and then assess the uncertainty based on the diversity of these samples (the LUQ metric). Green highlights indicate consistency across responses (low uncertainty) and red highlights discrepancies (high uncertainty). The LUQ-ENSEMBLE method selects the response from the LLM with the lowest uncertainty score as the final answer.

- We first highlight the limitations of existing UQ methods for long text generation and then propose LUQ (Long-text Uncertainty Quantification; pronounced as *luck*), a novel UQ method that computes sentence-level consistency in long text scenarios.
- Through extensive experiments on the original FACTSCORE dataset and our newly proposed FACTSCORE-DIS dataset in medical domain, we demonstrate that LUQ consistently shows strong negative correlations with the responses' factuality over 6 popular LLMs, outperforming all the baseline methods.
- We propose an ensemble modeling approach that selects responses from the model exhibiting the lowest LUQ uncertainty score, observing an improvement of up to 5% in the overall factuality scores. Additionally, we enhance the model's uncertainty awareness by implementing a selective answering strategy.

2 Background

2.1 Uncertainty and Confidence

Confidence and uncertainty in the context of machine learning models pertain to the level of assurance or certainty associated with a prediction or decision (Geng et al., 2023). While many studies treat *confidence* and *uncertainty* as antonyms and use them interchangeably (Xiao et al., 2022; Chen and Mueller, 2023), Lin et al. (2023) provide a clear distinction: uncertainty denotes the dispersion of potential predictions for a given input, whereas confidence pertains to the degree of confidence in a

specific prediction or output. We will adopt this terminology in the following sections.

Currently, a formal and universally accepted definition of uncertainty levels in language generation tasks remains elusive. Common practice in existing literature measures uncertainty through the entropy of predictions, akin to approaches in classification tasks (Kuhn et al., 2023; Lin et al., 2023). Predictive entropy is formally expressed as:

$$H(Y | x) = - \int p(y | x) \log(p(y | x)) dy$$

It captures the uncertainty associated with a prediction for a given input x . In the context of NLG, where \mathbf{R} denotes all possible generations and \mathbf{r} is a specific response, the uncertainty score can thus be conceptualized as:

$$U(x) = H(\mathbf{R} | x) = - \sum_{\mathbf{r}} p(\mathbf{r} | x) \log(p(\mathbf{r} | x))$$

In classification tasks, confidence for a specific prediction y is quantified using the predicted probability, represented as $\hat{p}(Y = y | x)$ (Geifman and El-Yaniv, 2017; Hendrycks and Gimpel, 2017). Similarly, in the context of NLG, the confidence score for a given response \mathbf{r} is represented by the joint probability of the tokens in the response:

$$C(x, \mathbf{r}) = \hat{p}(\mathbf{r} | x) = \prod_i \hat{p}(r_i | r_{<i}, x).$$

2.2 Uncertainty for Long Text Generation

In our study, we adopt a more flexible approach to defining uncertainty and confidence in long text

generation. Similar to Huang et al. (2024a), we focus on the ability of UQ methods to effectively rank responses, differentiating between correct and incorrect predictions. This approach also aligns with the concept of *relative confidence* as discussed by Geng et al. (2023). Our objective diverges from the orthogonal research direction about models’ calibration, which requires models to precisely reflect their true accuracy in practical scenarios (Lin et al., 2023). We argue that while short-answer questions may be straightforwardly assessed using metrics such as accuracy or exact match, these standards are often unrealistic for long text generation, given the complexities of real-life probabilities.

From a practical perspective, we aim for the uncertainty score to serve as a reliable indicator of the model’s performance. This performance encompasses several dimensions of generation quality, including factuality, coherence, and creativity. Our study prioritizes factuality and the truthfulness of responses, adopting these as our primary metrics. The factuality of the responses \mathbf{R} given a specific query x is denoted as $F(\mathbf{R} | \mathbf{x})$. Considering two inputs x_i and x_j , we explore the relationship between the model’s uncertainty, denoted as $U(\mathbf{x})$, and the factuality. Our goal is to have:

$$U(\mathbf{x}_i) \leq U(\mathbf{x}_j) \iff F(\mathbf{R} | \mathbf{x}_i) \geq F(\mathbf{R} | \mathbf{x}_j)$$

Correspondingly, for a given input x , the model’s confidence in generating a specific response r is represented as $C(\mathbf{x}, \mathbf{r})$. Thus, we aim to establish the following relationship:

$$C(\mathbf{x}, \mathbf{r}_i) \leq C(\mathbf{x}, \mathbf{r}_j) \iff F(\mathbf{r}_i | \mathbf{x}) \leq F(\mathbf{r}_j | \mathbf{x})$$

3 LUQ

In this section, we introduce our LUQ method and its two variations (LUQ-ATOMIC and LUQ-PAIR) to estimate uncertainty in long text generation. The overall framework is illustrated in Figure 1.

Motivation. Our underlying assumption posits that the greater the model’s uncertainty regarding a given question x , the more diverse its responses to question x will be. For instance, as shown in Figure 1, the term “*third pharaoh of the 20th Dynasty of Egypt*” is frequently supported by other sample responses, indicating the model’s high confidence in this information. However, the samples suggest different reign periods for Ramesses IV; the inconsistency shows the model’s higher uncertainty.

Following the generation of n responses, traditional UQ methods for short text commonly calculate the pairwise similarity among the responses (Kuhn et al., 2023; Lin et al., 2023). These pairwise similarity scores indicate the consistency between a pair of responses and play a vital role in subsequent uncertainty estimation. However, answers to certain questions such as “*Give me an introduction of ...*” and “*Tell me something about ...*” may extend to hundreds of words. Longer text leads to an unexpected high similarity across all response pairs when applying previous methods. To address this issue and achieve a more nuanced similarity assessment, we propose the LUQ uncertainty measurement with sentence-level similarity computation. Inspired by the hallucination detection method in Manakul et al. (2023), we split each response to sentences, and check whether each sentence can be supported by other samples.

Notation. Let r_a represent the response generated by a LLM to a user query x . We generate an additional n stochastic LLM sample responses $R = \{r_1, r_2, \dots, r_n\}$ using the same query. The set $R' = \{r_a, r_1, r_2, \dots, r_n\}$ encompasses all outputs from the model.

For any given response $r_i \in R'$, the first objective is to determine how often it is supported (or entailed) by other samples. To this end, we employ an NLI classifier to assess the similarity between r_i and each $r' \in R' \setminus \{r_i\}$. The output from an NLI classifier normally includes classifications of entailment, neutral, and contradiction, along with their respective logit values. It is important to note that we focus exclusively on the “entailment” and “contradiction” classes, as sentences labeled as “neutral” generally do not impact the overall factuality of a response. We calculate the NLI score for each sentence s_j within a response r , and then average these scores. Formally, the similarity score $\mathcal{S}(r_i, r')$ between r_i and r' is defined as:

$$\mathcal{P}(\text{entail} | s_j, r') = \frac{\exp(l_e)}{\exp(l_e) + \exp(l_c)}$$

$$\mathcal{S}(r_i, r') = \frac{1}{n} \sum_{j=1}^n \mathcal{P}(\text{entail} | s_j, r')$$

where l_e and l_c are the logits of the “entailment” and “contradiction” classes, respectively. We opt to calculate $\mathcal{P}(\text{entail} | s_j, r')$ over $\mathcal{P}(\text{contradict} | s_j, r')$ because non-contradictory responses can still be largely irrelevant, indicating higher uncertainty (Lin et al., 2023). The model’s confidence in

response r_i and the overall uncertainty is therefore defined as:

$$C(x, r_i) = \frac{1}{n} \sum_{r' \in R' \setminus \{r_i\}} \mathcal{S}(r_i, r')$$

$$U(x) = \frac{1}{n+1} \sum_{r_i \in R'} (1 - C(x, r_i))$$

Unlike [Kuhn et al. \(2023\)](#)’s method of applying an off-the-shelf DeBERTa model, we apply the DeBERTa-v3-large model ([He et al., 2023](#)), fine-tuned on the MultiNLI ([Williams et al., 2018](#)) dataset. This choice is due to our input being a concatenation of short hypothesis (sentence s) and a comparatively longer premises (reference response r'). The format of our input aligns with the task in MultiNLI dataset, ensuring an effective assessment of consistency among the responses.

LUQ-ATOMIC. To check the consistency of the generated responses in a more fine-grained manner, we implement LUQ-ATOMIC, a variation of the original LUQ. The key difference is that it first uses ChatGPT to break a response r into atomic fact pieces $\{a_1, a_2, \dots, a_j\}$. LUQ-ATOMIC then calculates the uncertainty scores based on atomic fact piece level (a_j) instead of sentence level (s_j).

LUQ-PAIR. The performance of our NLI classifier may be constrained by the length of the premises and hypotheses. To address this, we propose LUQ-PAIR to calculate the entailment score s_j for each sentence s'_j in r' and select the maximum value. Formally, we define this as:

$$\mathcal{P}(\text{entail} \mid s_j, r') = \max_{s'_j \in r'} |\mathcal{P}(\text{entail} \mid s_j, s'_j)|$$

We discuss more about LUQ-ATOMIC and LUQ-PAIR in [Appendix A](#).

4 Experiments

4.1 Dataset, Metric, and LLM Selection

Dataset. When selecting the dataset, we considered three main criteria: (1) The dataset should be a long-form QA dataset. (2) There should be a well-designed and widely-accepted automatic evaluation tool. (3) The questions should be clear, specific, and have definite answers for objective evaluation. We therefore employ FACTSCORE ([Min et al., 2023](#)) to evaluate the factuality of our generated text. It offers automated assessment with a low error rate (below 2%), enabling scalable application to diverse LLMs without requiring manual annotation. To supplement the extensive reliability

testing of FACTSCORE conducted by its creators, we performed a smaller-scale human annotation study. Our findings demonstrate a strong Pearson correlation of 0.88 between FACTSCORE ratings and human factuality judgments, confirming it being a reliable reference for factuality. Please refer to [Appendix B](#) for more information about the discussion of this dataset and our validation process.

The original FACTSCORE dataset (denoted as FACTSCORE-BIO) includes 500 individuals’ biographies from Wikidata with corresponding Wikipedia entries. To evaluate the applicability of UQ methods across different domains, we additionally developed a dataset, FACTSCORE-DIS, focusing on disease entities. Details of this dataset can be found in [Appendix C](#).

Metrics. For each generated response, FACTSCORE calculates a factuality score (FS). We apply FACTSCORE for the first generated response (r_a). As the LLMs may sometime refuse to answer certain questions, to have a fair comparison, we introduce a penalized factuality score (PFS) and penalized uncertainty score (PUS). To calculate PFS and PUS, we assign a factuality score of zero and uncertainty score of one to questions that models opt not to answer.

We then proceed to calculate both the Pearson Correlation Coefficient (PCC) and Spearman Correlation Coefficient (SCC) between the factuality scores and uncertainty scores. Following the criteria proposed by [Schober et al. \(2018\)](#), we classify the correlation coefficients into five categories based on their absolute values: over 0.9 indicates a very strong correlation; 0.7 to 0.9 signifies strong; 0.5 to 0.7 suggests moderate; 0.3 to 0.5 denotes weak; 0.1 to 0.3 implies very weak; and below 0.1 means negligible correlation.

LLMs. We selected six top-performing LLMs from the Arena Leaderboard ([Zheng et al., 2023](#)) for our experiments. Within our access rights, we chose three closed-sourced models: GPT-4 ([OpenAI, 2023](#)), GPT-3.5 ([OpenAI, 2022](#)), and Gemini 1.0 Pro ([Gemini Team, 2023](#)); and three open-sourced models: Yi-34B-Chat ([01.ai, 2023](#)), Tulu-2-70B ([Iverson et al., 2023](#)), and Vicuna-33B ([Zheng et al., 2023](#)). For each LLM, we include the following baseline UQ methods for comparison. Our implementation is based on the LM-Polygraph framework as proposed by [Fadeeva et al. \(2023\)](#). More details are provided in [Appendix D](#).

Baselines for UQ. We use the following black-

		White-Box Methods			Black-Box Methods						
		MSP	MCSE	SE	LexSim	Ecc	NumSets	EigV	Deg	SCN	LUQ
FACTSCORE-BIO											
GPT-4	PCC	-	-	-	-45.2	-24.8	-8.24	-36.9	-3.78	-53.1	-60.4
	SCC	-	-	-	-36.0	-12.7	4.18	-18.7	6.73	-41.8	-45.3
GPT-3.5	PCC	-	-	-	-67.8	-10.6	-11.9	-30.3	-22.4	-65.1	-71.3
	SCC	-	-	-	-52.4	-26.5	-17.0	-34.6	-22.9	-61.1	-66.6
Gemini 1.0 Pro	PCC	-	-	-	-67.2	-50.3	-53.0	-72.7	-64.4	-84.5	-85.1
	SCC	-	-	-	-63.7	-57.8	-57.0	-69.7	-67.7	-82.4	-81.3
Yi-34B-Chat	PCC	-20.7	-43.9	-55.8	-70.1	-27.6	-25.7	-49.0	-39.8	-70.3	-73.8
	SCC	-22.1	-44.3	-53.6	-68.2	-45.0	-31.3	-51.1	-38.9	-72.7	-74.6
Tulu-2-70B	PCC	-16.8	-32.4	-50.5	-55.7	-2.13	-20.7	-50.1	-53.4	-75.6	-77.6
	SCC	-15.4	-34.8	-52.7	-61.8	10.1	-18.1	-50.3	-54.0	-76.9	-75.4
Vicuna-33B	PCC	-28.5	-36.8	-58.6	-38.3	-18.7	-20.0	-60.5	-58.3	-66.8	-71.8
	SCC	-27.9	-37.4	-57.2	-50.6	-14.0	-16.6	-61.7	-62.4	-66.5	-70.8
FACTSCORE-DIS											
GPT-3.5	PCC	-	-	-	-41.8	-27.9	-7.81	-38.8	-13.5	-59.0	-67.3
	SCC	-	-	-	-39.4	-26.0	-6.94	-36.9	-16.3	-59.1	-65.3
Yi-34B-Chat	PCC	-20.3	-35.4	-52.6	-63.6	-19.3	-11.2	-40.6	-26.5	-65.1	-70.5
	SCC	-21.7	-33.8	-54.9	-58.7	-21.5	-16.3	-38.4	-22.1	-67.8	-72.4

Table 1: Pearson and Spearman correlation coefficients (expressed as percentages) between different LLMs and various UQ methods on the FactScore dataset. We use the original factuality scores instead of the penalized ones.

	FS	PFS	US	PUS	RR
GPT-4	80.8	72.4	20.8	29.0	86.6
GPT-3.5	68.3	68.3	25.7	25.7	100
Yi-34B-Chat	55.7	55.7	41.3	41.3	100
Tulu-2-70B	47.2	47.2	55.8	55.8	100
Gemini 1.0 Pro	43.2	42.7	61.7	62.2	98.9
Vicuna-33B	42.5	42.5	55.3	55.3	100

Table 2: Results on the FACTSCORE-BIO: FS and PFS are average and penalized factuality scores; US and PUS are average and penalized uncertainty scores by LUQ; RR is the response rate. All values are percentages.

box UQ methods as baselines: Lexical similarity (LexSim) (Fomicheva et al., 2020), Number of semantic sets (NumSets) (Lin et al., 2023), Sum of eigenvalues of the graph Laplacian (EigV) (Lin et al., 2023), Degree matrix (Deg) (Lin et al., 2023), Eccentricity (Ecc) (Lin et al., 2023), SelfCheckNLI (SCN) (Manakul et al., 2023). We also include three white-box methods for comparison: Maximum Sequence Probability (MSP), Monte Carlo Sequence Entropy (MCSE) (Malinin and Gales, 2021), and Semantic Entropy (SE) (Kuhn et al., 2023). More details can be found in Appendix E.

4.2 Uncertainty Quantification Results

Effectiveness of LUQ. Table 1 and Figure 2 illustrate the correlation between factuality scores and uncertainty scores. The results highlight LUQ’s effectiveness as an indicator of model factuality in long text generation tasks. LUQ demonstrates a strong negative correlation for GPT-3.5, Gemini 1.0

Pro, Yi-34B-Chat, Vicuna-33B, and Tulu-2-70B, with the strongest Pearson correlation being -0.851. For the baseline methods, LexSim emerges as a robust baseline offering lower computational demands. The confidence-based SCN method demonstrates the best Spearman correlation in models such as Gemini 1.0 Pro and Tulu-2-70B. Other baselines such as Ecc, NumSets and Deg yield unsatisfactory results, occasionally exhibiting even positive correlations. Case studies of our proposed LUQ method can be found in Appendix J.

Variations of LUQ. We compare the original LUQ with its two variations, LUQ-ATOMIC and LUQ-PAIR, in Table 5. We find that, with more fine-grained entailment checking, both consistently outperform the original LUQ. Further discussion on the pros and cons of these variations, along with usage guidelines, can be found in Appendix A.

LUQ for GPT-4. We also observe that LUQ is better suited for models with relatively lower factuality and a lack of self-expressiveness regarding uncertainty. For models with high factuality capabilities, such as GPT-4, LUQ only demonstrates a moderate correlation with factuality scores. As shown in Table 2, among all models, GPT-4 exhibits the highest overall factuality scores and the lowest average uncertainty scores. Figure 2a also shows that the data points of GPT-4 are tightly clustered with only few instances of high uncertainty. This is because GPT-4 tends to abstain from answering questions

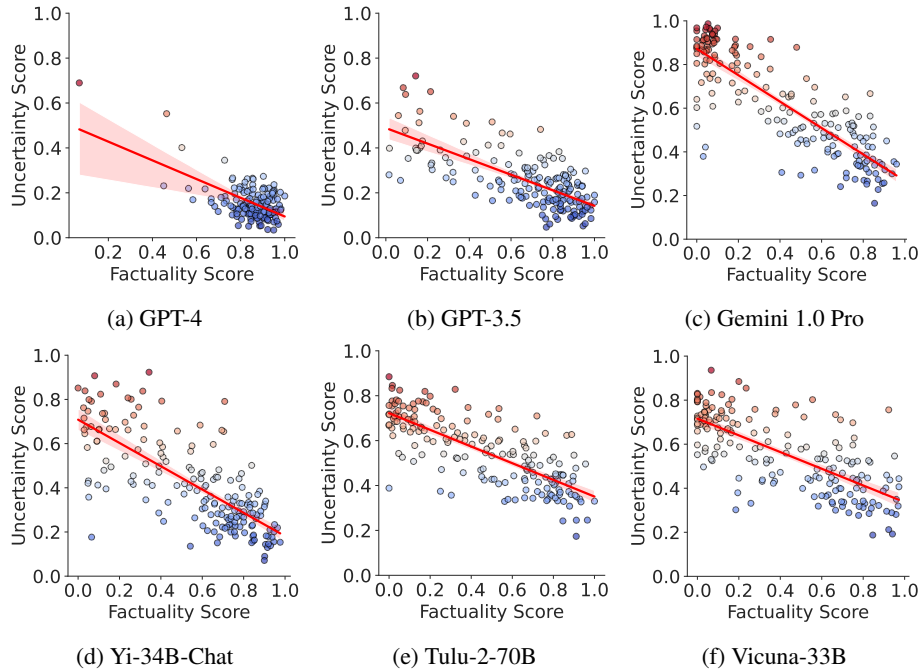


Figure 2: Scatter plot illustrating the relationship between factuality scores (x-axis) and uncertainty scores (y-axis) for different LLMs. Each point symbolizes an item in the FactScore dataset, with a red line highlighting the Pearson correlation. The distribution suggests a pattern where higher factuality correlates with lower uncertainty.

more often compared to other models, highlighting improved uncertainty self-detection. However, this observation does not influence the effectiveness of our method, as in real life models with lower factuality and unable to express uncertainty are in greater need of external uncertainty measurements.

LUQ in FACTSCORE-DIS. We test one closed-source LLM, GPT-3.5, and one open-source LLM, Yi-34B-Chat in our newly proposed FACTSCORE-DIS. Our LUQ model consistently surpasses the performance of baseline models, thereby demonstrating its effectiveness on the newly proposed dataset within the medical domain.

Higher frequency leads to higher factuality and lower uncertainty. In Figure 3, we compare the factuality and uncertainty scores across different entity frequencies. The original FACTSCORE dataset provides the frequency of each entity in Wikipedia, categorizing them based on page views and co-occurrence within the training set (Min et al., 2023). Frequencies are classified into five categories, ranging from “very rare” to “very frequent.” Our observations suggest that questions associated with higher entity frequencies tend to yield more factual responses, alongside decreased model uncertainty. Notably, GPT-4 demonstrates consistent performance regarding uncertainty and factuality across varying frequencies, potentially attributable

to its selective response strategy. Although it answers all the questions in the “very frequent,” “frequent,” and “medium” categories, it refuses to answer around 25% of “rare” questions and 30% of “very rare” questions.

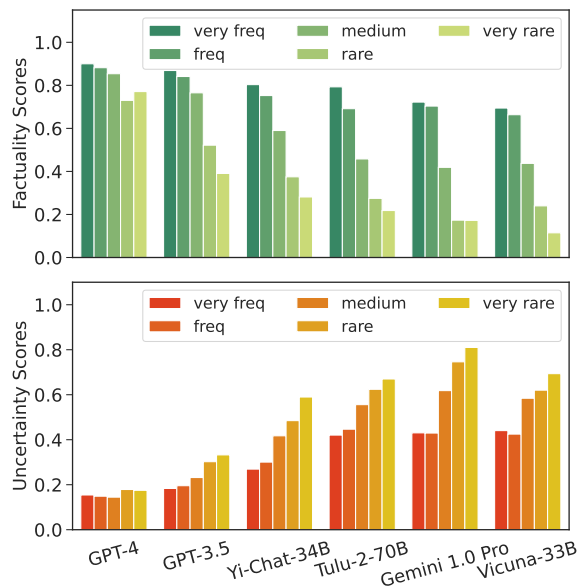


Figure 3: Factuality and uncertainty scores across different frequencies on FACTSCORE-BIO.

4.3 LUQ-ENSEMBLE

Given the variance in training corpus, different LLMs may possess varying levels of knowledge for

Percentile	GPT-3.5		Yi-34B-Chat		Tulu-2-70B		Vicuna-33B		Gemini 1.0 Pro	
	FS	US	FS	US	FS	US	FS	US	FS	US
0	68.3	25.7	55.7	41.3	47.2	55.8	42.5	55.3	43.2	61.7
2.5	69.8	24.1	56.9	40.2	48.3	53.9	43.6	54.4	44.3	60.0
5	70.8	23.4	58.0	39.3	49.4	53.1	44.5	53.8	45.2	59.2
7.5	71.5	22.7	58.9	38.1	50.3	52.6	45.5	53.0	46.3	58.2
10	72.3	22.2	60.2	36.8	51.4	51.9	46.1	52.3	47.3	57.7
12.5	74.1	21.6	61.7	35.0	52.1	51.3	46.5	51.6	48.4	56.3
15	75.0	21.2	62.9	34.2	53.3	50.6	47.5	51.0	49.5	55.4

Table 3: Selective question answering results on FACTSCORE-BIO (expressed as percentage). The percentile indicates the percentage of questions for which answers were abstained.

Methods	PFS	PUS	AD
Tulu-2-70B	47.2	55.8	42.1
Gemini 1.0 Pro	42.7	62.2	29.5
Vicuna-33B	42.5	58.1	28.4
LUQ-ENSEMBLE	52.8	45.8	100
Yi-34B-Chat	55.7	41.3	66.1
Tulu-2-70B	47.2	55.8	21.3
Gemini 1.0 Pro	42.7	62.2	12.6
LUQ-ENSEMBLE	58.8	37.6	100
GPT-3.5	67.3	25.7	92.4
Gemini 1.0 Pro	42.7	62.2	1.64
Vicuna-33B	42.5	58.1	6.01
LUQ-ENSEMBLE	67.4	24.8	100
GPT-4	72.1	29.0	60.1
GPT-3.5	67.3	25.7	32.8
Yi-34B-Chat	55.7	41.3	7.10
LUQ-ENSEMBLE	76.6	17.3	100

Table 4: Results of different ensemble strategies on FACTSCORE-BIO (expressed as percentage). The Answer Distribution (AD) indicates the percentage of final answers generated by each component model.

a specific question. After obtaining outputs from different LLMs, the challenge now is to *choose the best one without fact-checking each answer* (which is both time-consuming and costly (Guo et al., 2022; Zhang et al., 2024)). Utilizing the LUQ uncertainty score as a reliable indicator of factuality, we enhance overall performance through an ensemble approach. In this method, the model exhibiting the lowest LUQ score for a given question is chosen as the final answer. Experimental results (Table 4) affirm the superiority of the LUQ-ENSEMBLE over its constituent counterparts.

Ensembling models with similar factuality scores can notably enhance performance. Our findings suggest that ensembling models with similar factuality scores can significantly enhance performance. For instance, in the combination of Tulu-2-70B, Gemini 1.0 Pro, and Vicuna-33B, the PFS increases by 5% compared to the originally

top-performing Tulu-2-70B, which scored 47.19%. Additionally, ensembling models with comparable performance leads to a more balanced distribution of answers. In contrast, integrating a model with substantially superior performance, as seen in the combination of GPT-3.5, Gemini 1.0 Pro, and Vicuna-33B, predominantly favors answers from GPT-3.5 (92.35%), leading to marginal improvement (0.06%) in the ensemble method.

Ensembling does not guarantee better performance. While ensembling always reduces uncertainty scores (as we select the model with the least uncertainty), it does not necessarily improve factuality scores. Ensembling LLMs according to poor UQ methods may result in overall performance that is worse than that of its individual components. Table 10 in Appendix G compares the effectiveness of using LUQ as the ensemble indicator with other methods. *The results indicate that ensembling does not inherently enhance performance.* For example, with UQ method Ecc, the ensemble factuality score can be lower than that of its best-performing component (47.2% vs 43.3%). In contrast, using LUQ as the ensembling indicator yields the best overall performance.

4.4 Selective Question Answering

From Table 2, it is observed that while GPT-4 opts not to respond to some queries, other models generally attempt to answer all questions. The limited refusal by Gemini 1.0 Pro primarily stems from considerations of sensitive content and regulatory constraints, rather than uncertainty. Therefore, we investigate the application of the LUQ score to equip these models with the capability for selective question answering—that is, to enable them to decline responses when uncertain. Contrary to the traditional aim of responding correctly to every question, the objective in a selective question answering framework is to preserve accuracy while

maximizing the number of questions answered (Kamath et al., 2020; Cole et al., 2023; Yang et al., 2023; Dong et al., 2024).

Table 3 presents the results of selective question answering. The models are permitted to refrain from answering questions with high uncertainty. The *percentiles* indicate the proportion of questions each model abstained from answering. The findings demonstrate that adopting a selective answering approach enhances the models’ factuality by allowing for more question rejections. By declining to answer a similar proportion of questions (approximately 15%) as GPT-4, the models typically achieve an improvement of over 5% in overall factuality scores. In Appendix H, we provide a detailed discussion on how practitioners can use LUQ to implement selective answering strategies, including setting and adjusting uncertainty thresholds.

5 Related Work

UQ in Machine Learning. Prior to LLMs, UQ has been extensively explored within the field of machine learning (Gawlikowski et al., 2023). According to the source of uncertainty, it is typically categorized into two types: aleatoric and epistemic uncertainty (Hora, 1996; Der Kiureghian and Ditlevsen, 2009). Aleatoric uncertainty, also known as statistical uncertainty, pertains to the inherent randomness in experimental outcomes due to stochastic effects (Hüllermeier and Waegeman, 2021). In contrast, epistemic uncertainty stems from incomplete knowledge, potentially including uncertainties in a machine learning model’s parameters or the lack of certain training data (Hüllermeier and Waegeman, 2021; Huang et al., 2023). Our focus is primarily on epistemic uncertainty.

UQ in LLMs. In contrast to discriminative models, which readily provide probability scores for specific categories, uncertainty estimation in generative LLMs presents unique challenges: (1) There is an exponential increase in the output space as sentence length grows, rendering the evaluation of all possible predictions impractical (Geng et al., 2023; Wang et al., 2023). (2) The significance of semantic nuances and their inherent uncertainties, which diverges from the fixed category labels typical of discriminative models, complicates matters further (Kuhn et al., 2023). Generally, UQ methods for LLMs can be categorized based on the accessibility of the model’s internal states, distinguishing between black-box and white-box approaches.

White-box LLMs often rely on logit-based evaluations, assessing sentence uncertainty through token-level probabilities or entropy (Murray and Chiang, 2018; Kuhn et al., 2023; Vazhentsev et al., 2023; Duan et al., 2023).

However, as access to LLMs increasingly relies on API calls, research has pivoted towards black-box methods. These can be further categorized into: (i) *verbalized methods*, which prompt LLMs to articulate their uncertainty in the output, using phrases like “I am sure” or “I do not know” (Mielke et al., 2022). Nonetheless, a practical mismatch between the expressed and actual uncertainty levels has been noted (Lin et al., 2022; Xiong et al., 2023). Xiong et al. (2023) highlight that LLMs often display excessive confidence when verbalizing their certainty. (ii) *Consistency-based (sampling-based)* estimation premises on the assumption that increased uncertainty in a model corresponds to greater diversity in its outputs, frequently resulting in hallucinatory outputs (Manakul et al., 2023; Lin et al., 2023). Our proposed method, LUQ, follows this consistency-based approach. There are also efforts on integrating verbalized methods with consistency-based approaches (Xiong et al., 2023; Rivera et al., 2024). Understanding uncertainty in LLMs can enhance in-context learning (Zhou et al., 2023; Li et al., 2023), selective question answering (Yang et al., 2023), LLM cascading (Huang et al., 2024b), adaptive retrieval (Ding et al., 2024), language agents (Han et al., 2024), and model self-refinement (Yao et al., 2024; Chen et al., 2024).

6 Conclusion

In this work, we first identify that existing UQ methods are ineffective on long text generation. We therefore introduce LUQ, a novel UQ method tailored for long-form text generation in LLMs. It overcomes the limitation of previous methods by calculating sentence level consistency. We conduct extensive experiments over six popular LLMs, such as GPT-4 and Gemini 1.0 Pro. We extend the existing FACTSCORE dataset with human validation and annotations for additional disease domain. Our findings demonstrate that LUQ significantly improves the correlation with models’ factuality scores over previous methods across various different setups and domains. LUQ serves as a reliable indicator of model’s factuality performance. Additionally, we present LUQ-ENSEMBLE, a model ensembling and selective question answering strat-

egy, which showcases a promising avenue for enhancing the factual accuracy of LLM outputs. This research not only advances our understanding of UQ in the context of LLMs but also offers practical tools for improving the reliability and trustworthiness of AI-generated content.

Limitation

The limitations of this study include the following: (1) A primary challenge in studying uncertainty quantification for long text generation lies in the difficulty of evaluating the generated text. Unlike classification tasks and short-answer QA, there is no straightforward metric for assessing the quality of generated text. In this study, we employ the factuality score as the primary evaluation metric, thereby leaving other text aspects, such as coherence, cohesion, and creativity, under-explored. Future work could investigate uncertainty scores using more comprehensive evaluation metrics. (2) In this study, we do not investigate the performance of UQ methods under ambiguous and unanswerable questions, such as ASQA (Stelmakh et al., 2022) and SelfAware (Yin et al., 2023). Previous uncertainty metrics for short-answer questions are tested on answerable questions with clear intentions. This is because clearly defined questions with definite answers provide a straightforward framework for evaluating model accuracy. In contrast, unanswerable or ambiguous questions lack clear ground truths, complicating the assessment of uncertainty estimates. We advocate researchers to explore more in this area in the future. (3) We focus on assessing the overall uncertainty of a model, rather than model uncertainty on individual instances. The relative uncertainty equation in Section 2.2 represents an ideal scenario. If a model learns a significant amount of non-factual data over factual data for a particular entity/instance, the aforementioned equation can be inaccurate for that case. Future work could investigate the causes of this special case and develop strategies to address it during the pre-training stage.

Ethics Statement

Our research adheres to rigorous ethical guidelines, with a strong emphasis on data privacy, bias mitigation, and societal impact. During the dataset construction phase, we verified the licenses of all software utilized and ensured strict compliance with these licenses. We have thoroughly assessed our

project and do not foresee any other potential risks.

Acknowledgments

Caiqi Zhang is supported by an Amazon Studentship. We thank Chengzu Li, Yulong Chen, and the reviewers for their valuable feedback on this paper.

References

- 01.ai. 2023. Building the next generation of open-source and bilingual llms. <https://huggingface.co/01-ai/Yi-34B-Chat>. Accessed: 2024-02-05.
- Anthropic. 2023. Introducing claude 2.1. Available from Anthropic: <https://www.anthropic.com/news/claude-2-1>.
- Joris Baan, Nico Daheim, Evgenia Ilia, Dennis Ulmer, Haau-Sing Li, Raquel Fernández, Barbara Plank, Rico Sennrich, Chrysoula Zerva, and Wilker Aziz. 2023. *Uncertainty in natural language generation: From theory to applications*.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2023. *A survey on evaluation of large language models*. *ArXiv preprint*, abs/2307.03109.
- Jiuhai Chen and Jonas Mueller. 2023. *Quantifying uncertainty in answers from any language model and enhancing their trustworthiness*.
- Yulong Chen, Yang Liu, Jianhao Yan, Xuefeng Bai, Ming Zhong, Yinghao Yang, Ziyi Yang, Chenguang Zhu, and Yue Zhang. 2024. *See what llms cannot answer: A self-challenge framework for uncovering llm weaknesses*.
- Jeremy Cole, Michael Zhang, Daniel Gillick, Julian Eisenschlos, Bhuwan Dhingra, and Jacob Eisenstein. 2023. *Selectively answering ambiguous questions*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 530–543, Singapore. Association for Computational Linguistics.
- Armen Der Kiureghian and Ove Ditlevsen. 2009. *Aleatory or epistemic? does it matter?* *Structural safety*, 31(2):105–112.
- Hanxing Ding, Liang Pang, Zihao Wei, Huawei Shen, and Xueqi Cheng. 2024. *Retrieve only when it needs: Adaptive retrieval augmentation for hallucination mitigation in large language models*.
- Yijiang River Dong, Tiancheng Hu, and Nigel Collier. 2024. *Can llm be a personalized judge?*
- Jinhao Duan, Hao Cheng, Shiqi Wang, Chenan Wang, Alex Zavalny, Renjing Xu, Bhavya Kailkhura, and

- Kaidi Xu. 2023. [Shifting attention to relevance: Towards the uncertainty estimation of large language models](#). *ArXiv preprint*, abs/2307.01379.
- Ekaterina Fadeeva, Roman Vashurin, Akim Tsvigun, Artem Vazhentsev, Sergey Petrankov, Kirill Fedyanin, Daniil Vasilev, Elizaveta Goncharova, Alexander Panchenko, Maxim Panov, Timothy Baldwin, and Artem Shelmanov. 2023. [LM-polygraph: Uncertainty estimation for language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 446–461, Singapore. Association for Computational Linguistics.
- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [ELI5: Long form question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3558–3567, Florence, Italy. Association for Computational Linguistics.
- Marina Fomicheva, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Francisco Guzmán, Mark Fishel, Nikolaos Aletras, Vishrav Chaudhary, and Lucia Specia. 2020. [Unsupervised quality estimation for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:539–555.
- Jakob Gawlikowski, Cedrique Rovile Njiteucheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. 2023. [A survey of uncertainty in deep neural networks](#). *Artificial Intelligence Review*, 56(Suppl 1):1513–1589.
- Yonatan Geifman and Ran El-Yaniv. 2017. [Selective classification for deep neural networks](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 4878–4887.
- Gemini Team. 2023. [Gemini: A family of highly capable multimodal models](#). Technical report, Google.
- Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepl, Preslav Nakov, and Iryna Gurevych. 2023. [A survey of language model confidence estimation and calibration](#). *ArXiv preprint*, abs/2311.08298.
- Zhijiang Guo, Michael Schlichtkrull, and Andreas Vlachos. 2022. [A survey on automated fact-checking](#). *Transactions of the Association for Computational Linguistics*, 10:178–206.
- Jiuzhou Han, Wray Buntine, and Ehsan Shareghi. 2024. [Towards uncertainty-aware language agent](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 6662–6685, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2023. [Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Dan Hendrycks and Kevin Gimpel. 2017. [A baseline for detecting misclassified and out-of-distribution examples in neural networks](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Stephen C Hora. 1996. [Aleatory and epistemic uncertainty in probability elicitation with an example from hazardous waste management](#). *Reliability Engineering & System Safety*, 54(2-3):217–223.
- Mengting Hu, Zhen Zhang, Shiwan Zhao, Minlie Huang, and Bingzhe Wu. 2023. [Uncertainty in natural language processing: Sources, quantification, and applications](#).
- Xinmeng Huang, Shuo Li, Mengxin Yu, Matteo Sesia, Hamed Hassani, Insup Lee, Osbert Bastani, and Edgar Dobriban. 2024a. [Uncertainty in language models: Assessment through rank-calibration](#).
- Yuheng Huang, Jiayang Song, Zhijie Wang, Huaming Chen, and Lei Ma. 2023. [Look before you leap: An exploratory study of uncertainty measurement for large language models](#). *ArXiv preprint*, abs/2307.10236.
- Yukun Huang, Yixin Liu, Raghuveer Thirukovalluru, Arman Cohan, and Bhuwan Dhingra. 2024b. [Calibrating long-form generations from large language models](#).
- Eyke Hüllermeier and Willem Waegeman. 2021. [Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods](#). *Machine Learning*, 110:457–506.
- Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. 2023. [Camels in a changing climate: Enhancing lm adaptation with tulu 2](#).
- Amita Kamath, Robin Jia, and Percy Liang. 2020. [Selective question answering under domain shift](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5684–5696, Online. Association for Computational Linguistics.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. [Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.

- Chengzu Li, Han Zhou, Goran Glavaš, Anna Korhonen, and Ivan Vulić. 2023. [On task performance and model calibration with supervised and self-ensembled in-context learning](#).
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [Teaching models to express their uncertainty in words](#). *ArXiv preprint*, abs/2205.14334.
- Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. 2023. [Generating with confidence: Uncertainty quantification for black-box large language models](#).
- Andrey Malinin and Mark J. F. Gales. 2021. [Uncertainty estimation in autoregressive structured prediction](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Potsawee Manakul, Adian Liusie, and Mark Gales. 2023. [SelfCheckGPT: Zero-resource black-box hallucination detection for generative large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9004–9017, Singapore. Association for Computational Linguistics.
- Sabrina J. Mielke, Arthur Szlam, Emily Dinan, and Y-Lan Boureau. 2022. [Reducing conversational agents’ overconfidence through linguistic calibration](#). *Transactions of the Association for Computational Linguistics*, 10:857–872.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FActScore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Singapore. Association for Computational Linguistics.
- Kenton Murray and David Chiang. 2018. [Correcting length bias in neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 212–223, Brussels, Belgium. Association for Computational Linguistics.
- OpenAI. 2022. [Chatgpt blog post](#).
- OpenAI. 2023. [Gpt-4 technical report](#).
- Mauricio Rivera, Jean-François Godbout, Reihaneh Rabbany, and Kellin Pelrine. 2024. [Combining confidence elicitation and sample-based methods for uncertainty quantification in misinformation mitigation](#). In *Proceedings of the 1st Workshop on Uncertainty-Aware NLP (UncertainLP 2024)*, pages 114–126, St Julians, Malta. Association for Computational Linguistics.
- Patrick Schober, Christa Boer, and Lothar A Schwarte. 2018. [Correlation coefficients: appropriate use and interpretation](#). *Anesthesia & analgesia*, 126(5):1763–1768.
- Ivan Stelmakh, Yi Luan, Bhuwan Dhingra, and Ming-Wei Chang. 2022. [ASQA: Factoid questions meet long-form answers](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8273–8288, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Artem Vazhentsev, Akim Tsvigun, Roman Vashurin, Sergey Petrakov, Daniil Vasilev, Maxim Panov, Alexander Panchenko, and Artem Shelmanov. 2023. [Efficient out-of-domain detection for sequence to sequence models](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1430–1454, Toronto, Canada. Association for Computational Linguistics.
- Cunxiang Wang, Xiaoze Liu, Yuanhao Yue, Xiangru Tang, Tianhang Zhang, Cheng Jiayang, Yunzhi Yao, Wenyang Gao, Xuming Hu, Zehan Qi, Yidong Wang, Linyi Yang, Jindong Wang, Xing Xie, Zheng Zhang, and Yue Zhang. 2023. [Survey on factuality in large language models: Knowledge, retrieval and domain-specificity](#).
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Yuxin Xiao, Paul Pu Liang, Umang Bhatt, Willie Neiswanger, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2022. [Uncertainty quantification with pre-trained language models: A large-scale empirical analysis](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 7273–7284, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. 2023. [Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms](#).
- Qi Yang, Shreya Ravikumar, Fynn Schmitt-Ulms, Satvik Lolla, Ege Demir, Iaroslav Elistratov, Alex Lavaee, Sadhana Lolla, Elaheh Ahmadi, Daniela Rus, Alexander Amini, and Alejandro Perez. 2023. [Uncertainty-aware language modeling for selective question answering](#).
- Yuxuan Yao, Han Wu, Zhijiang Guo, Biyan Zhou, Jiahui Gao, Sichun Luo, Hanxu Hou, Xiaojin Fu, and Linqi Song. 2024. [Learning from correctness without prompting makes llm efficient reasoner](#).
- Zhangyue Yin, Qiushi Sun, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Xuanjing Huang. 2023. [Do large language models know what they don’t know?](#) In

Findings of the Association for Computational Linguistics: ACL 2023, pages 8653–8665, Toronto, Canada. Association for Computational Linguistics.

Caiqi Zhang, Zhijiang Guo, and Andreas Vlachos. 2024. [Do we need language-specific fact-checking models? the case of chinese.](#)

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT.](#) In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023. [Siren’s song in the ai ocean: A survey on hallucination in large language models.](#) *ArXiv preprint*, abs/2309.01219.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models.](#) *ArXiv preprint*, abs/2303.18223.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhonghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena.](#) In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.

Han Zhou, Xingchen Wan, Lev Prolee, Diana Mincu, Jilin Chen, Katherine Heller, and Subhrajit Roy. 2023. [Batch calibration: Rethinking calibration for in-context learning and prompt engineering.](#)

A LUQ variations

Table 5 compares LUQ, LUQ-PAIR, and LUQ-ATOMIC. Our findings indicate that LUQ-PAIR and LUQ-ATOMIC outperform the original LUQ across all models.

The superiority of LUQ-PAIR stems from its use of shorter premises for NLI (sentence s'_j instead of r'), which leads to higher NLI accuracy. However, this improvement comes at the cost of increased computational requirements. For N samples with M sentences each, the original LUQ requires $M \times M$ NLI computations, whereas LUQ-PAIR requires $N \times M^2$ computations.

In LUQ-ATOMIC, we first break down the text into atomic sentences using ChatGPT before proceeding with further steps. The main concern of

this variation is about **evaluation fairness**. Both LUQ-ATOMIC and FACTSCORE use ChatGPT to break long texts into atomic sentences, potentially creating an unfair comparison with other UQ methods that do not involve this step. Further thorough investigation is needed to determine if this approach is universally beneficial, regardless of the atomic fact producer/converter used. This would require a new study and could be a valuable follow-up work. Notably, our original LUQ can *still outperform existing baselines without this step*.

Regarding the choice of LUQ and its variations, we recommend the following:

1. For general purposes and scenarios requiring high efficiency, use LUQ.
2. For cases needing very accurate uncertainty estimation and where time is not a constraint, use LUQ-ATOMIC if the budget allows for API calls. If not, use LUQ-PAIR.

B Dataset Selection

B.1 Dataset for Long-form Uncertainty

As mentioned in Section 4.1, when selecting the dataset, we considered three main criteria: (1) The dataset should be a long-form QA dataset with relatively lengthy answers. (2) There should be a well-designed and widely-accepted automatic evaluation tool. (3) The questions should be clear, specific, and have definite answers for objective evaluation.

Evaluating long-form QA is a long-standing challenge, making the last criterion especially important to mitigate factors that could affect evaluation quality. According to Hu et al. (2023) and Baan et al. (2023), uncertainty in NLG systems can be disentangled into three main sources: *input*, *model*, and *output*. Due to the intrinsic ambiguity of language and unknown queries, the input itself contains uncertainty (Baan et al., 2023). To conduct a **more controlled study**, we focus primarily on *output uncertainty*, assuming all questions are generally answerable and clearly stated.

Among all the datasets, FACTSCORE advances the field by using LLMs for human-level evaluation, addressing the limitations of traditional metrics like BLEU, ROUGE-L, and BERTScore. Other long-form QA benchmarks fall short in at least one criterion. For example, ELI5 (Fan et al., 2019) questions are very general (e.g., “How can different animals perceive different colors?”) and can be answered in many ways, making it hard to define

	GPT-4	GPT-3.5	Yi-34B-Chat	Tulu-2-70B	Gemini 1.0 Pro	Vicuna-33B
LUQ	-60.4	-71.3	-73.8	-77.6	-85.1	-71.8
LUQ-PAIR	-61.3	-72.8	-76.1	-80.5	-86.1	-72.9
LUQ-ATOMIC	-63.5	-75.6	-79.6	-83.6	-87.2	-75.7

Table 5: Pearson Correlation Scores between factuality scores and uncertainty scores for LUQ’s variations on FACTSCORE-BIO dataset.

objective criteria for a good answer. ASQA (Stelmakh et al., 2022) questions are inherent ambiguous, making it unsuitable for proving the efficiency of an UQ method.

B.2 Human Evaluation on FACTSCORE

We also engaged human annotators to assess the factuality of the generated passages. Although Min et al. (2023) conducted comprehensive experiments to demonstrate the effectiveness of the FACTSCORE framework, we perform a sanity check by directly correlating the annotated passage factuality with uncertainty scores. We recruited three students with Master’s degrees in Computer Science from our university to conduct the human annotations. We ensured the annotators were not involved in our project and had not discussed it. We used Fleiss’ Kappa to measure inter-annotator agreement, achieving a score of 0.793, indicating substantial agreement (close to the “almost perfect” standard of 0.8-1.0) according to Landis and Koch (1977). Annotators are compensated above the local minimum hourly wage standard. The instructions provided to the annotators are listed in Figure 6.

We randomly selected 50 passages from the responses generated by the Yi-34B-Chat model. We observed a Pearson correlation coefficient of 0.88 between the FACTSCORE factuality score and the human-annotated factuality score. This finding aligns with the results reported by Min et al. (2023), demonstrating that FACTSCORE is a reliable tool in our experiments. Table 7 compares the results of different UQ methods with those obtained using FACTSCORE and human annotation.

C FACTSCORE-DIS

To demonstrate the generalization of our proposed LUQ model across various domains, we create a new dataset adopting the methodology used to construct the original FACTSCORE dataset for the disease entities. To differentiate, we refer the original dataset as FACTSCORE-BIO and the new dataset

as FACTSCORE-DIS. The detailed information of FACTSCORE-DIS dataset is as follows:

Data Collection Following FACTSCORE-Bio, we use Wikipedia as our main knowledge source. We first select all the diseases names using the following SPARQL codes calling the wiki API. We then removed those diseases with empty Wikipedia pages.

Following FACTSCORE-BIO, we utilized Wikipedia as our primary knowledge source. Initially, we extracted all disease names using the following SPARQL queries to call the Wikidata API. Subsequently, we removed those diseases with empty Wikipedia pages.

```
SELECT ?item ?itemLabel WHERE {
  ?item wdt:P31 wd:Q112193867. # is an
  instance of class of diseases
  SERVICE wikibase:label { bd:
    serviceParam wikibase:language "[
  AUTO_LANGUAGE],en". }
```

Frequency For each entity retrieved, we adhere to the methodology described by Min et al. (2023) to assign a frequency label ranging from “Very Rare” to “Very Frequent” based on an entity’s pageviews. It’s crucial to acknowledge that in the context of diseases, the number of diagnosed cases is commonly used as a metric. However, we opted not to use this metric because our goal is to simulate the distribution of these diseases within the training corpus of LLMs. Relying solely on diagnosed case numbers may underrepresent the prominence of a disease within the corpus. Diseases like Amyotrophic Lateral Sclerosis (ALS), despite their low incidence rate in the population, attract significant global interest and impact. As a result, LLMs may demonstrate extensive knowledge about such diseases, reflecting their visibility in the data on which they are trained, rather than their actual morbidity rates.

After determining the frequencies, we sampled 36 disease entities for each category, amassing a total of 180 data points. Subsequently, we conducted a human evaluation to validate the selected

Human Annotation Guidelines

Your task is to evaluate the veracity of each sentence in the provided passage. It is crucial to carefully assess each statement for accuracy and relevance to the main topic.

Steps to Follow:

1. **Read the Passage Thoroughly:** Begin by reading the entire passage to grasp the overall context and the main topic.
2. **Check Each Sentence:** Examine each sentence individually for accuracy and completeness. Determine if the information is factual and supported by reliable sources, and whether the sentence presents a partial truth or is fully accurate.
3. **Scoring:** Assign each sentence a score based on its accuracy, using a specified range (e.g., 1 to 3). Scores should reflect:
 - The sentence is entirely accurate and provides a complete picture. [Highest Score: 3]
 - The sentence is partially correct but may lack context or omit important details. [Mid-Range Score: 2]
 - The sentence is largely inaccurate or misleading. [Lowest Score: 1]
4. **Relevance:** Flag any sentence that does not contribute to or is off-topic as *Not Relevant*.

Guidelines:

- Use reliable sources (e.g. Wikipedia) to verify factual information, maintaining an impartial stance throughout.
- Keep the passage and your assessments confidential.

Table 6: Human Annotation Guidelines

Methods	FactScore		Human	
	PCC	SCC	PCC	SCC
LexSimilarity	-67.3	-66.4	-65.6	-64.0
Eccentricity	-26.3	-25.5	-22.6	-25.1
NumSemSets	-26.4	-26.9	-24.3	-23.5
EigValLaplacian	-45.8	-43.9	-43.4	-42.7
DegMat	-38.9	-39.7	-36.8	-31.6
SelfCheckNLI	-68.5	-67.3	-66.1	-69.2
LUQ	-72.7	-71.4	-69.0	-68.3

Table 7: Pearson and Spearman correlation coefficients (expressed as percentages) between different factuality scores and various UQ methods on the **FactScore-Bio** dataset using Yi-34B-Chat.

data points, replacing any that were deemed unsuitable with diseases that were more clearly defined and well-documented. Several examples from the dataset are showcased in Table 8.

D Experiment Setup

For GPT-4 and GPT-3.5, we use the OpenAI API, with specific version `gpt-4-turbo-0125-preview` and `gpt-3.5-turbo-0613`. For Gemini 1.0 Pro, we call the API for developers. For Yi-34B-Chat, Tulu-2-70B (`tulu-2-dpo-70b`), and Vicuna-33B (`vicuna-33b-v1.3`), we use them off-the-shelf and only for inference. The temperature is set to 0.7. We run our uncertainty measurement experiments on A100-SXM-80GB GPUs. For our experiments, we use the following prompt:

Tell me a short bio of the person <entity>. Begin with their birth, significant life events, achievements, and contributions. Include their education, career milestones, any notable awards or recognitions received, and their impact on their field or society. Ensure the biography is concise, factual, and engaging, covering key aspects of their life and work.

From the estimation of [Min et al. \(2023\)](#), run-

Frequency	Wikidata ID	Disease Name
Very Freq	Q8071861	Zika fever
	Q12199	HIV/AIDS
	Q12152	myocardial infarction
	Q12206	diabetes
	Q12204	tuberculosis
Freq	Q154874	yellow fever
	Q188638	mood disorder
	Q159701	glaucoma
	Q1138580	Ewing sarcoma
	Q209369	Hodgkin lymphoma
Medium	Q5134736	cloacal exstrophy
	Q247978	anisometropia
	Q2373361	tree nut allergy
	Q778731	pyuria
	Q7900433	urethral syndrome
Rare	Q220322	agnosia
	Q2735907	cutis laxa
	Q500695	retinoblastoma
	Q627625	histoplasmosis
	Q1347729	Epstein syndrome
Very Rare	Q21505502	spina bifida
	Q1862031	pinguecula
	Q1361850	patulous eustachian tube
	Q4667534	leiomyoma
	Q595010	hypertrichosis

Table 8: Frequency Categories of Diseases

ning FACTSCORE costs about \$1 of the API cost per 100 sentences. For instance, for 100 generations, each with 5 sentences on average, it costs \$5 in total.

E Baselines

We mainly use the library LM-Polygraph (Fadeeva et al., 2023) for the UQ methods. Here we provide a brief introduction for each method:

LexicalSimilarity (Fomicheva et al., 2020): it computes the similarity between two phrases using metrics like ROUGE scores and BLEU. For our experiment, we utilize BERTScore (Zhang et al., 2020) to enhance performance, computing the average similarity score with other answers.

NumSemSets (Lin et al., 2023): it clusters semantically equivalent answers into the same sets. Initially, the number of semantic sets equals the total number of generated answers. Then it sequentially examines responses, making pairwise comparisons between them, and combines different answers. One of the limitation of this method is that the uncertainty score $U_{NumSemSets}$ can only take inte-

ger values. EigValLaplacian is therefore designed to overcome this problem.

EigValLaplacian (Lin et al., 2023): For a similarity matrix S , it calculates the Normalized Graph Laplacian of S using $L = I - D^{-\frac{1}{2}}SD^{-\frac{1}{2}}$, where D is a diagonal matrix and $D_{ii} = \sum_{j=1}^m S_{ij}$ (m is the number of responses). Consequently, the uncertainty score is defined as $U_{EigV} = \sum_{k=1}^m \max(0, 1 - \lambda_k)$. This value is a continuous analogue of $U_{NumSemSets}$. In extreme case if adjacency matrix S is binary these two measures will coincide.

DegMat (Lin et al., 2023): it is based on the idea that the total uncertainty of the answers might be measured as a corrected trace of the diagonal matrix D . This is because elements on the diagonal of matrix D are sums of similarities between the given answer and other answers. We thus define uncertainty estimate $U_{Deg}(x) = \text{trace}(m - D)/m^2$.

Eccentricity (Lin et al., 2023): A drawback of previously considered methods is the limited knowledge of the actual embedding space for the different answers since we only have measures of their similarities. The graph Laplacian, however, can provide us with coordinates for the responses. Denote $\mathbf{u}_1, \dots, \mathbf{u}_k \in \mathbb{R}^m$ as the eigenvectors of L that correspond to k smallest eigenvalues. We can efficiently construct an informative embedding $\mathbf{v}_j = [\mathbf{u}_{1,j}, \dots, \mathbf{u}_{k,j}]$ for an answer \mathbf{y}_j . Then it uses the average distance from center as the uncertainty measure, defined as: $U_{Ecc} = \|\frac{1}{m} \sum_{\ell=1}^m \tilde{\mathbf{v}}_{\ell}^T\|_2$, where $\tilde{\mathbf{v}}_j = \mathbf{v}_j - \frac{1}{m} \sum_{\ell=1}^m \mathbf{v}_{\ell}$.

SelfCheckNLI (Manakul et al., 2023): As defined in Section 2, SelfCheckNLI primarily functions as a confidence measurement tool, calculating the similarity exclusively between the primary response r_a and the other generated samples. Distinctively, it evaluates $\mathcal{P}(\text{contradict} | s, r')$ and focuses solely on $C(x, r_a)$.

F Response Statistics

Average Response Length. In our study, we define ‘‘long text’’ as the model’s output consisting of at least 100 words. For our experiments, the texts are even longer, averaging more than 200 words. The average word count is calculated and shown in the following table. In contrast, commonly used datasets for existing UQ methods have much shorter texts, with an average of 1.95 words for Trivia QA dataset and 3.37 words for Natural Questions dataset.

Model	Avg. Word Count per Response
GPT-4-0125	264
GPT-3.5-turbo-0613	239
Gemini Pro 1.0	223
Yi-34B-Chat	307
Tulu-2-70B	267
Vicuna-33B	206

Table 9: The average response length for each LLM.

Number of Facts in a Response Figure 4 shows the average atomic facts provided by various AI models for the FACTSCORE dataset. GPT-4 has the highest average number of atomic facts at 52.24, indicating it provides the most detailed factual responses. Tulu-2-70B follows with an average of 52.17, nearly matching GPT-4 in factual details. GPT-3.5 has an AF of 50.67, showing it also delivers a high level of factual details in its responses. Yi-34B-Chat and Gemini 1.0 Pro have comparatively lower averages, at 45.80 and 42.72 respectively. Vicuna-33B has the lowest AF at 36.20, indicating it offers the least amount of factual information in its responses. Generally, these models provide similar number of atomic facts in their responses.

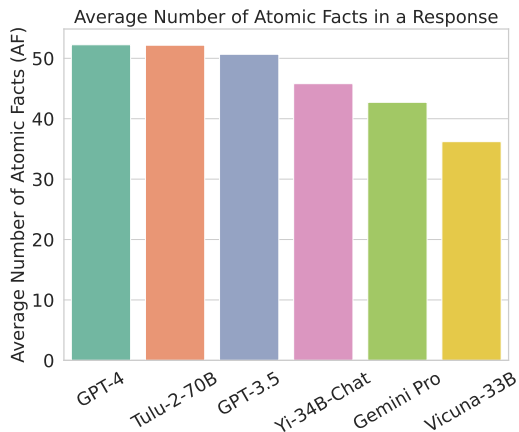


Figure 4: Average number of atomic Facts (AF) in a response for each model.

G LUQ-ENSEMBLE

In this section, we discuss more about the novelty, motivation, and effectiveness of our LUQ-ENSEMBLE method.

The Motivation of LUQ-ENSEMBLE. After getting multiple answers from different LLMs, the challenge is now to *choose which one as the final output*. Traditional aggregation methods like majority vote and weighted vote are ineffective for

long text generation because finding the majority answer is difficult when all the responses are somewhat different, and methods like boosting or bagging require additional training. Our uncertainty measure thus serves as an effective indicator for model ensembling.

The Effectiveness of LUQ-ENSEMBLE. Meanwhile, it is not guaranteed that ensembling will lead to better performance. It is true that the uncertainty scores will by definition decrease (as we select the model with the least uncertainty as the final response), but the factuality score may not. The effectiveness of LUQ-ENSEMBLE largely relies on the reliability of the UQ method.

Table 10 compares the effectiveness of using LUQ as the ensemble indicator with other methods. The results show that ensembling *does not inherently improve performance*. With a poor UQ method (such as Ecc), the ensemble factuality score can be lower than that of its highest component (47.2 vs 43.3). In contrast, using LUQ as the indicator for ensembling yields the best overall performance.

	PFS	PUS
Individual Results		
Tulu-2-70B	47.2	55.8
Gemini 1.0 Pro	42.7	62.2
Vicuna-33B	42.5	58.1
Different Ensemble Methods		
ECC-ENSEMBLE	43.4	35.5
LEXSIM-ENSEMBLE	47.6	39.8
SELF-CHECK-ENSEMBLE	49.3	46.7
LUQ-ENSEMBLE	52.8	45.8

Table 10: Penalised factuality score (PFS) and Penalised uncertainty score (PUS) for individual models and ensembles with different UQ methods.

H Selective QA Strategy

When implementing a selective answering strategy in practical applications, it is essential for practitioners to tailor the uncertainty thresholds to the specific models and tasks at hand. In our experiment, as shown in Table 2, we find that GPT-4 tends to refuse to answer around 15% of the questions. To simulate a GPT-4-like answering strategy, for each model, we set different thresholds to ensure they refuse to answer between 0 and 15% of the questions. Our experiments indicate that different LLMs may have varying average absolute

uncertainty values, making a universal uncertainty threshold unsuitable for all models. Additionally, the inherent nature of the tasks may influence practitioners’ decisions to make the model more conservative or more willing to attempt answering users’ questions.

We advise practitioners to implement selective QA strategies using the following practical steps:

- Collect a representative set of questions/queries that closely mimic real-world usage scenarios.
- Obtain responses from the LLMs for these questions and apply UQ methods (e.g., LUQ) to get uncertainty scores.
- Establish thresholds tailored to the specific task and the practitioners’ goals, selecting either cautious (lower threshold) or lenient (higher threshold) settings as needed.
- Develop clear strategies for handling high uncertainty, such as refusing to answer, requesting clarification, or using alternative approaches.

I Ablation Study

Temperature As the diversity of content generated by LLMs may be influenced by the temperature setting, we adjust the temperature to test the robustness of our methods. Due to limitations in computational resources and API budget constraints, we selected GPT-3.5, Yi-34B-Chat, and Vicuna-33B for our experiments (refer to Figure 5). Our findings indicate that a lower temperature leads to a weaker correlation score, likely because the generated responses are more uniform, providing limited information for the self-consistency test. As the temperature increases, we observe a strengthening in correlation. However, beyond a certain point, further increases in temperature lead to diminishing improvements and can even result in a weaker correlation. We hypothesize that excessively diverse responses may complicate the NLI process, as a greater number of sentences fail to be supported by other samples.

Number of Samples Previous research on short answer generation (Kuhn et al., 2023; Lin et al., 2023) has demonstrated that an increase in the number of samples correlates with enhanced performance. We investigate whether it also applies

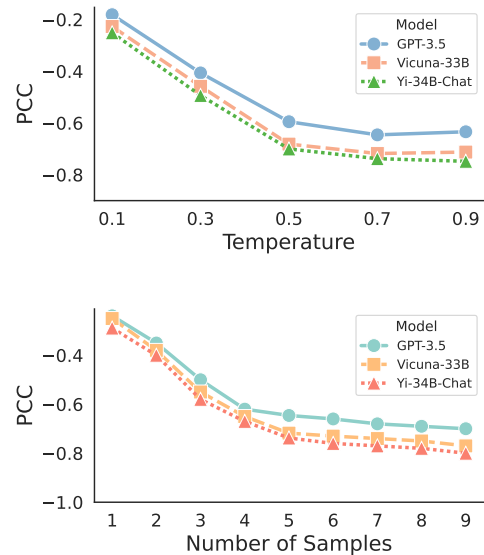


Figure 5: The effect of different temperatures (upper) and the number of samples (lower) on the PCC with LUQ.

to long-text generation and find that with more samples, LUQ shows better performance and PCC scores, which corroborates with previous observations in short-text generation, as depicted in Figure 5. Providing a greater number of samples enables the NLI process to predict sentence factuality with higher accuracy. However, a notable drawback of increasing the sample size is the associated rise in computational costs.

J Case Study

In this section, we present two case studies illustrating the performance of LUQ. For simplicity, we only show three samples for each question using Yi-34B-Chat. *In both cases, the baseline UQ methods indicate low uncertainty.* However, the first case is with higher factuality, while the second case includes many non-factual information. Unlike baseline UQ methods that can only calculate similarity at the passage level, our method captures more nuanced differences.

In the first case study, LUQ demonstrates lower uncertainty. As shown in Figure 6, we use green to represent sentences that can be cross-validated by other samples. It is evident that most sentences are consistent across different samples, thus leading to lower uncertainty.

In the second case study, we showcase an example with higher LUQ uncertainty. In Figure 7, red is used to indicate contradictory information. It is apparent that the samples are not consistent with each other. The samples even state different nationalities of Jonathan Haagensen.



Tell me a bio of Chadwick Boseman .

Chadwick Boseman was an American actor, playwright, and screenwriter who rose to prominence for his portrayal of historical figures, most notably as T'Challa in the Marvel Cinematic Universe films "Black Panther," "Captain America: Civil War," "Avengers: Infinity War," and "Avengers: Endgame." Boseman was born on November 29, 1976, in Anderson, South Carolina. He graduated from Howard University and began his acting career in theater, appearing in various stage productions. His first major film role was as baseball player Jackie Robinson in the biopic "42" (2013). He went on to play other notable figures, such as James Brown in "Get on Up" (2014) and Thurgood Marshall in "Marshall" (2017). Boseman's portrayal of T'Challa, the Black Panther, in the Marvel Cinematic Universe, beginning with "Captain America: Civil War" in 2016, brought him widespread recognition and critical acclaim. "Black Panther," which was released in 2018, was a cultural phenomenon, breaking box office records and earning him a NAACP Image Award and a Screen Actors Guild Award for his performance. Off-screen, Boseman was known for his philanthropy and commitment to education, often visiting schools and speaking about the importance of representation in media. He was also an advocate for cancer awareness, having battled colon cancer for several years, which he kept private. Sadly, Chadwick Boseman passed away on August 28, 2020, at the age of 43, due to complications from colon cancer. His death was a great loss to the entertainment industry and his fans around the world. His legacy lives on through his powerful performances and the impact he had on diversity and representation in film.

Chadwick Boseman was an American actor, playwright, and screenwriter who rose to prominence for his portrayals of real-life historical figures, most notably Jackie Robinson in "42" (2013) and James Brown in "Get on Up" (2014). He was born on November 29, 1977, in Anderson, South Carolina. Boseman graduated from Howard University with a degree in Fine Arts in 2000. After a few years of performing in theater and making appearances in television series, he landed his breakthrough role as Jackie Robinson in "42," which earned him critical acclaim. He followed this with a star-turning performance as James Brown in "Get on Up," further establishing himself as a talent to watch. Boseman's career reached new heights when he was cast as T'Challa, the Black Panther, in the Marvel Cinematic Universe, beginning with "Captain America: Civil War" (2016). The role catapulted him to international stardom, and he reprised the character in "Black Panther" (2018), which he also executive produced. "Black Panther" was a cultural phenomenon, breaking box office records and becoming the first superhero film to be nominated for the Academy Award for Best Picture. He was also an advocate for education and often spoke about the importance of literacy and education in his interviews. Sadly, Boseman passed away on August 28, 2020, at the age of 43, following a private four-year battle with colon cancer. His death sent shockwaves through the entertainment industry and beyond, with tributes pouring in from all over the world. His legacy as an artist and a role model continues to inspire, and his impact on cinema and culture is profound.

Chadwick Boseman was an American actor, playwright, and filmmaker, best known for his groundbreaking role as T'Challa in the Marvel Cinematic Universe film "Black Panther." Born on November 29, 1977, in Anderson, South Carolina, Boseman graduated from Howard University and began his acting career in theater, appearing in various stage productions. Boseman's breakthrough came with his portrayal of legendary baseball player Jackie Robinson in the 2013 film "42." He went on to play other notable figures, such as James Brown in "Get on Up" (2014) and Thurgood Marshall in "Marshall" (2017). However, it was his role as the Black Panther that catapulted him to international fame. As T'Challa, Boseman became an inspiration to many, representing a powerful and complex superhero who was also the king of the fictional African nation of Wakanda. "Black Panther" was not only a critical and commercial success but also a cultural phenomenon, breaking barriers and inspiring a new generation of superhero fans. Sadly, Boseman passed away on August 28, 2020, at the age of 43, after a private four-year battle with colon cancer. His death was a shock to the world, and his legacy as an artist and a role model continues to be celebrated and honored. Chadwick Boseman's life and work have left an indelible mark on the entertainment industry and society at large.

Figure 6: Example with **low uncertainty** by LUQ. Sentences that can be cross-validated by other samples are shown in green.



Tell me a bio of Jonathan Haagensen.

Jonathan Haagensen is a Danish actor, model, and singer who was born on October 27, 1979, in Copenhagen, Denmark. He is best known for his role as Marcus Brostrup in the Danish television series "Borgen," which aired from 2010 to 2013. Haagensen began his career as a model, working for various fashion brands and appearing in magazines. He later transitioned into acting, landing his breakthrough role in "Borgen," where he played the role of the Prime Minister's press secretary. His performance in the series was well-received, and it helped to establish him as a prominent figure in the Danish entertainment industry. Aside from "Borgen," Haagensen has appeared in several other Danish television shows and films. He has also ventured into music, releasing his debut album in 2014. In addition to his work in entertainment, Haagensen is also involved in various philanthropic activities. He has supported organizations that work towards environmental conservation and has been an ambassador for causes related to children's rights and education. Jonathan Haagensen continues to act, model, and pursue his musical career while maintaining a low profile personal life.

Jonathan Haagensen is a Brazilian actor, model, and television personality. He was born on January 28, 1983, in Rio de Janeiro, Brazil. Haagensen began his career as a model, working for various fashion brands and appearing in campaigns for companies such as Calvin Klein and Versace. He transitioned into acting, making his debut in the Brazilian television series "Malhação" in 2003. Haagensen gained widespread recognition for his role in the telenovela "Páginas da Vida" (2006), where he portrayed the character of Henrique. This role earned him a nomination for the Best Actor award at the 2007 International Emmy Awards. Haagensen has since appeared in several other Brazilian telenovelas, including "Caminho das Índias" (2009), "A Favorita" (2011), and "Amor à Vida" (2013). He has also ventured into film, with roles in movies like "O Cheiro do Ralo" (2006) and "As Cariocas" (2010). Apart from his acting career, Haagensen has also hosted various television shows and events. He is known for his good looks and has been considered one of the most handsome Brazilian actors. Haagensen continues to work in the entertainment industry, both in Brazil and internationally.

Jonathan Haagensen is a Brazilian actor, model, and television host. He was born on June 25, 1974, in Rio de Janeiro, Brazil. Haagensen began his career as a model and quickly gained popularity in the fashion industry, working for various international brands and gracing the covers of several magazines. In addition to his modeling career, Haagensen ventured into acting, making his debut in the telenovela "Coração de Estudante" in 2002. He has since appeared in several Brazilian television series and films, establishing himself as a prominent figure in the country's entertainment industry. Haagensen is also known for his work as a television host. He has hosted various shows, including the Brazilian version of the reality show "Big Brother." Throughout his career, Haagensen has been recognized with several awards and nominations for his performances in television and film. He continues to be an influential figure in Brazilian media, known for his talent and good looks.

Figure 7: Example with **high uncertainty** by LUQ. Contradictory information is highlighted in red.