

LoGU: Long-form Generation with Uncertainty Expressions

Ruihan Yang^{*,†}, Caiqi Zhang^{*,†}, Zhisong Zhang[†],
Xinting Huang[†], Sen Yang^{*}, Nigel Collier[†], Dong Yu[†], Deqing Yang[†]

^{*}Fudan University, [†]University of Cambridge, [‡]Tencent AI Lab,

[‡]The Chinese University of Hong Kong

{rhyang17, deqingyang}@fudan.edu.cn, cz391@cam.ac.uk, zhisongzhang@tencent.com

Abstract

While Large Language Models (LLMs) demonstrate impressive capabilities, they still struggle with hallucinations. A promising approach to mitigate hallucinations is enabling models to express uncertainty when unsure. Previous research on uncertainty estimation has primarily focused on short-form QA, but real-world applications often require much longer responses. In this work, we introduce the task of **Long-form Generation with Uncertainty (LoGU)**, which requires the models to explicitly express uncertainty during the generation. We identify two key challenges: *Uncertainty Suppression*, where models hesitate to express uncertainty, and *Uncertainty Misalignment*, where models convey uncertainty inaccurately. To tackle these challenges, we propose a novel decomposition-based data collection framework and a two-stage training pipeline. Specifically, we use supervised fine-tuning (SFT) for uncertainty suppression problem and direct preference optimization (DPO) for uncertainty misalignment. Experiments on three long-form datasets demonstrate the effectiveness of our approach, showing improvements in factual accuracy, reduction of incorrect statements, and preservation of the overall comprehensiveness of the generated responses. Further analysis reveals that baseline methods tend to express uncertainty in vague and broad terms, while our method generates more specific and targeted uncertainty expressions.¹

1 Introduction

While large language models (LLMs) demonstrate remarkable performance across various domains, they still suffer from a significant limitation: the generation of factually incorrect statements (*i.e.*, hallucinations) (Zhao et al., 2023; Hadi et al., 2023; Chang et al., 2024; Chen et al., 2024). This

^{*}Equal contribution. Work done during the internship at Tencent AI Lab. [†]Corresponding authors.

¹Project page: <https://github.com/rhyang2021/LoGU>.

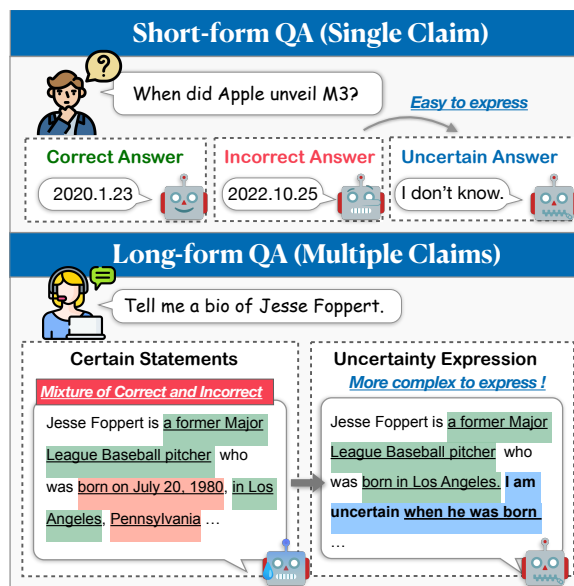


Figure 1: Existing works on short-form QA only give an overall uncertainty to the query. On the contrary, the long-form QA responses consist of both factually correct and incorrect claims, requiring more fine-grained uncertainty expressions.

issue hinders the broader adoption of LLMs in real-world applications that demand highly reliable and accurate responses (Zhang et al., 2024a,b; Jiang et al., 2025). Alleviating hallucinations remains a challenging and long-standing problem (Gekhman et al., 2024; Lin et al., 2024). LLMs often produce responses even in the absence of sufficient knowledge (Xiong et al., 2024a; Kang et al., 2024). Recognizing these intrinsic limitations, **it is essential for LLMs to explicitly express uncertainty during generation.**

Previous work on enabling LLMs to express uncertainty has primarily focused on short-form question answering (QA) tasks (Lin et al., 2022; Tian et al., 2023; Li et al., 2024; Zhang et al., 2024c). However, real-world applications often require much longer responses, sometimes spanning hundreds or even thousands of words (Bai

et al., 2023; Zhang et al., 2024b; Yuan et al., 2024). As shown in Figure 1, unlike short-form responses, which typically contain a single claim, long-form responses involve multiple claims, increasing the complexity of accurately modeling uncertainty.

In this paper, we introduce the task of **Long-form Generation with Uncertainty (LOGU)**, which focuses on enabling LLMs to express uncertainty accurately in long-form responses. We emphasize linguistic expression of uncertainty rather than post-hoc estimation (Zhang et al., 2024b; Fadeeva et al., 2023). We argue that embedding uncertainty expressions directly in the generation is more human-interpretable and aligns more naturally with everyday communication.

We first identify two sub-challenges for LOGU: *Uncertainty Suppression* and *Uncertainty Misalignment*. *Uncertainty Suppression* refers to the reluctance of LLMs to express uncertainty, while *Uncertainty Misalignment* denotes cases where LLMs express uncertainty inaccurately. Ideally, LLMs should express uncertainty only when dealing with unknown facts. To address these challenges, we propose a novel decomposition-based data collection framework and a two-stage training pipeline.

As illustrated in Figure 2, during data collection, we first decompose responses into atomic claims, then selectively revise them to include uncertainty expressions, and finally reassemble the atomic claims. We use the collected data in a two-stage training pipeline consisting of LOGU-SFT and LOGU-DPO. In LOGU-SFT, we encourage the model to express uncertainty during generation to mitigate the *uncertainty suppression* problem. In LOGU-DPO, response pairs with accurate and inaccurate uncertainty expressions are optimized using Direct Preference Optimization (Rafailov et al., 2023) to address the *uncertainty misalignment*.

Experiments on three long-form datasets demonstrate the effectiveness of our approach, showing improvements in factual accuracy, reduction of incorrect statements, and preservation of the overall comprehensiveness of the generated responses. Further analysis reveals that baseline methods tend to express uncertainty in vague and broad terms, while our method generates more specific and targeted uncertainty expressions.

2 Task definition

Definition. We define the **Long-form Generation with Uncertainty (LOGU)** task as follows. A lan-

guage model, denoted by \mathcal{M} , is prompted with an open-ended query (e.g., "Introduce [a person] to me.") to generate a long-form response $R \sim \mathcal{M}(R | q)$ (e.g., a biography). The response R typically lacks explicit uncertainty expressions and includes both factually correct and incorrect statements. For fine-grained processing, we decompose the response R into N atomic claims C , represented as $C = \prod_{i=1}^N c_i$. Each atomic claim is self-contained with only one piece of information. These claims are categorized into three subsets: C_s , containing claims supported by external knowledge; C_{ns} , containing claims refuted by external knowledge; and C_{unc} , including claims with uncertainty expressions. The LOGU task requires the model to reduce C_{ns} and express uncertainty in C_{unc} more accurately.

Two Sub-Challenges. The LOGU task presents two main challenges: 1) *Uncertainty Suppression*: The model tends to avoid expressing uncertainty, displaying an overconfidence bias, which can result in reduced factual accuracy (Xiong et al., 2024b; Tian et al., 2023; Zhang et al., 2024b); 2) *Uncertainty Misalignment*: The model’s expressed uncertainty does not always align with its actual knowledge. It may show uncertainty even when it can provide correct answers.

Primary Goals. Ideally, a model excelling in LOGU should maximize the following two aspects:

- **Factual Accuracy (FA)** It represents the proportion of correctly generated deterministic claims in the model’s output, reflecting how often the model generates factually correct statements. A higher FA score indicates that the model is generating a larger proportion of correct facts relative to incorrect ones. Specifically, FA is defined as:

$$FA = \frac{|C_s|}{|C_s| + |C_{ns}|}.$$

- **Uncertain Accuracy (UA)** It represents the quality of uncertainty expressions by calculating how often the model accurately express uncertainty. Ideally, the model should express uncertainty in two specific cases: (1) When the model does not have sufficient knowledge about a fact or is prone to make errors; (2) When the uncertainty is expressed about specific, granular details (e.g., "I do not know [a person]’s birthday."), rather than general or vague statements (e.g., "I am not sure about [a person]’s early life.") is over

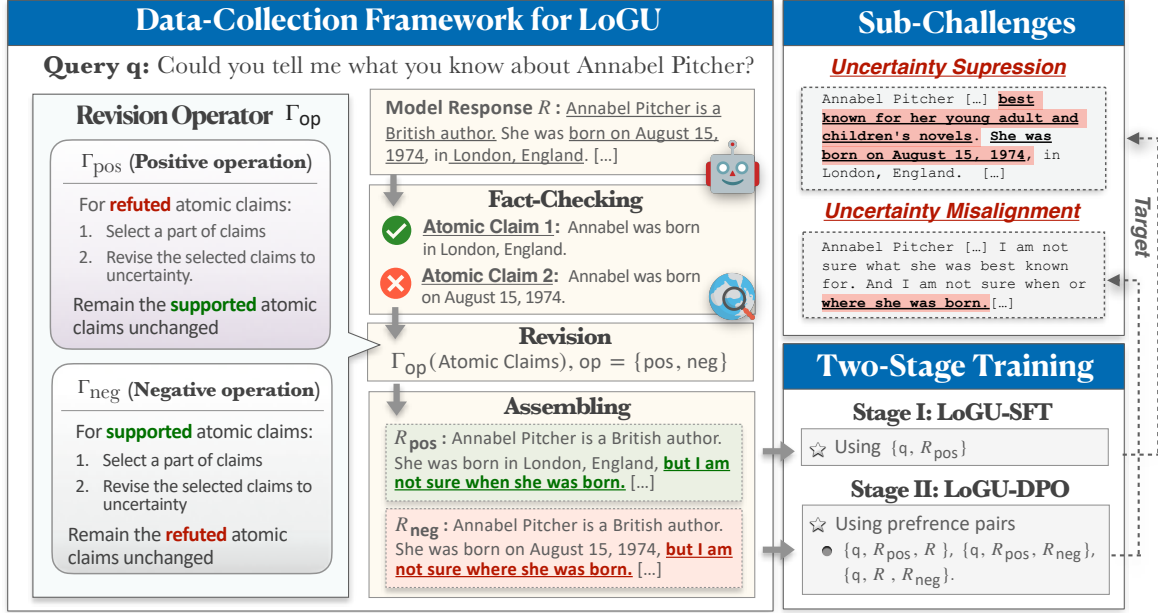


Figure 2: Our framework employs a decomposition-based approach: first, decompose responses into atomic claims for fact-checking and then revise selected atomic claims with uncertainty expressions. The revised atomic claims are assembled to a coherent response. The collected data is then used in the two-stage training pipeline: LOGU-SFT and LOGU-DPO, which address *uncertainty suppression* and *uncertainty misalignment*, respectively.

broad). Within the uncertainty expression set C_{unc} , we denote the truly uncertain claims as C_{unc}^{true} , and UA is defined as:

$$UA = \frac{|C_{unc}^{true}|}{|C_{unc}|}.$$

More discussion on FA and UA are in Section 4.2.

3 Methodology

In this section, we explain our approach for the LoGU task in details. Figure 2 shows an overview of the data collection and our training pipeline.

3.1 Data Collection Framework for LoGU

For a target LLM \mathcal{M} and a list of queries $\{q_i\}_{i=1}^n$, we first gather an original set of query-response pairs, denoted as $\mathcal{D} = \{(q_1, R_1), \dots, (q_n, R_n)\}$. Then we apply the following procedures to refine the responses to properly express uncertainty, with the help of an additional auxiliary LLM \mathcal{A} (e.g., GPT-4 (Achiam et al., 2023); See Appendix A for the prompts of the procedures).

Fact-Checking. We use the FACTSCORE (Min et al., 2023) for fact-checking. For each query q , the original response R is decomposed into a series of atomic claims C , with each claim containing only a single piece of information. This decomposition is performed by the auxiliary model \mathcal{A} .

Next, we retrieve relevant passages from external sources (e.g., Wikipedia or Google search) for each claim. The veracity of each claim is then assessed by model \mathcal{A} based on its entailment by these passages. Claims that are supported are classified as C_s , while those refuted are classified as C_{ns} .

Revision. We apply the operator Γ_{op} to the sets C_s and C_{ns} , where op specifies the type of revision: either positive (pos) or negative (neg). Specifically, Γ_{op} selects a subset of claims based on a given ratio and rewrite the selected claims into ones with uncertainty expressions. The details of Γ_{op} are further explained in §3.2.

Assembling. After revising the atomic claims extracted from the response R , we assemble these altered claims into a coherent and fluent refined response by prompting the auxiliary model \mathcal{A} . For the revised atomic claims $\Gamma_{op}(C_s)$ and $\Gamma_{op}(C_{ns})$, the assembled response is denoted as R_{op} . The final training dataset is represented as $\mathcal{D}_{op} = \{(q_1, R_{op,1}), \dots, (q_n, R_{op,n})\}$.

3.2 Revision Operator Γ_{op}

Two types of operator Γ_{op} (i.e., pos and neg) are used to construct training data to address two specific problems: *uncertainty suppression* and *uncertainty misalignment*.

Positive Revision Operation. To address *uncer-*

tainty suppression, we construct QA pairs with explicitly expressed uncertainty. We first define Γ_{pos} , which revises refuted atomic claims (C_{ns}) into uncertainty expressions, while leaving supported claims (C_{s}) unchanged. We rewrite the atomic claims by prompting the auxiliary model \mathcal{A} with ten manually designed distinct uncertainty expression patterns. For instance, the statement “[a person] was born on August 15, 1974” can be revised to “I am not sure when [a person] was born.”

In real-world applications, the proportion of uncertainty expressions in long-form responses can influence the overall user experience. To prevent an overuse of uncertainty expressions from reducing readability, we apply a down-sampling mechanism to control their ratio. We therefore selectively revise a subset of refuted claims from C_{ns} with a target uncertainty ratio α (more details about the impact of the uncertainty ratio in Appendix E). Specifically, we select $\min\left(\frac{\alpha}{1-\alpha}|C_{\text{s}}|, |C_{\text{ns}}|\right)$ claims from C_{ns} , ensuring the final responses meet the desired uncertainty ratio α . We denote the revised claims as $C_{\text{ns-unc}}$.

Negative Revision Operation. To address *uncertainty misalignment*, where the model might express uncertainty about information it actually knows, we construct negative learning instances that contain such inaccurate uncertainty expressions. We define the operator Γ_{neg} to revise supported atomic claims C_{s} into uncertainty expressions while leaving refuted claims C_{ns} unchanged. The output of Γ_{neg} will be used as penalization signals to let the model not to inappropriately express uncertainty. Similar to Γ_{pos} , we also control the uncertainty ratio at α . We select $\min\left(\frac{\alpha}{1-\alpha}|C_{\text{ns}}|, |C_{\text{s}}|\right)$ claims from C_{s} , which are then revised into uncertainty expressions, denoted as $C_{\text{s-unc}}$.

Summary. With the operations of Γ_{pos} and Γ_{neg} , we can rewrite an original response R into R_{pos} and R_{neg} , respectively. R_{pos} consists of the union of C_{s} and $C_{\text{ns-unc}}$, while R_{neg} includes $C_{\text{s-unc}}$ and C_{ns} . Note that since the down-sampling selection of atomic claims within C_{s} and C_{ns} is random, the revisions performed by Γ_{op} on C_{s} and C_{ns} are not unique. This results in multiple possible variants, from which we select up to five samples per instance for our training process.

3.3 Training Pipeline

With the constructed training data, we propose a two-stage training pipeline. In Stage 1 (LOGU-SFT), our goal is to enable the model to express uncertainty, specifically addressing the problem of uncertainty suppression. In Stage 2 (LOGU-DPO), we aim to enhance the model’s precision in expressing uncertainty, ensuring it avoids both exaggeration and understatement.

LOGU-SFT. We fine-tune the language model \mathcal{M} using only positively refined responses (R_{pos}). Specifically, the model iteratively processes a sequence t_1, t_2, \dots, t_T consisting of queries and corresponding responses (q, R_{pos}). The objective is to minimize the cross-entropy loss \mathcal{L} :

$$\mathcal{L} = -\frac{1}{T} \sum_{i=1}^T \log P(t_i | t_1, t_2, \dots, t_{i-1}),$$

where $P(t_i | t_1, t_2, \dots, t_{i-1})$ denotes the probability of predicting token t_i given all previous tokens. Importantly, the loss is computed exclusively for the response parts of the sequence, omitting the queries. After fine-tuning, we denote the model policy as π_{sft} .

LOGU-DPO. To refine the model’s alignment in expressing uncertainty, we construct response pairs of good and unsatisfactory expressions from triplets² of ($R_{\text{pos}}, R, R_{\text{neg}}$). The negative refinements (R_{neg}) are used to address uncertainty misalignment by providing negative examples of inappropriate uncertainty expressions.

With these constructed response pairs (q, y_w, y_l), we employ the Direct Preference Optimization (DPO) algorithm (Rafailov et al., 2023) for model training:

$$\mathcal{L}_{\theta} = -\mathbb{E}_{(q, y_w, y_l) \sim \mathcal{D}} \left[\log \sigma \left(\beta \log \frac{\pi_{\theta}(y_w | q)}{\pi_{\text{sft}}(y_w | q)} - \beta \log \frac{\pi_{\theta}(y_l | q)}{\pi_{\text{sft}}(y_l | q)} \right) \right],$$

where π_{θ} is the model policy initialized from π_{sft} , β controls the deviation from π_{sft} , and σ represents the logistic function.

4 Experiment Settings

4.1 Dataset and Models

We use three datasets for long-form QA: (1) *Bios* (Min et al., 2023), which contains 500 individuals

² R_{pos} is preferred over R , which is in turn preferred over R_{neg} . This results in $\binom{3}{2}$ preference pairs per query.

from Wikipedia with varying levels of popularity. The models are asked to generate biographies for the individuals. (2) *LongFact* (Wei et al., 2024), which extends *Bios* and includes 1,140 questions covering 38 manually-selected topics. (3) *Wild-Hallu* (Zhao et al., 2024) comprises 7,917 entities derived from one million user-chatbot interactions in real-world settings. We also conduct experiments on an out-of-domain dataset ASQA (Stelmakh et al., 2022) in § 5.1, to validate the generalization performance of our method. Regarding the models, we use the Llama3-8B-Instruct (Meta, 2024) and Mistral-7B-Instruct (Jiang et al., 2023). For auxiliary model \mathcal{A} , we use GPT-4o. Implementation details are in Appendix B.

4.2 Evaluation Metrics

We evaluate the performance of models in reflecting uncertainty in long-form generation using three core metrics: Factual Accuracy (FA), Uncertain Accuracy (UA), and number of incorrect claims (#Incor). As discussed in §2, the first two metrics are designed to measure the model’s ability to accurately output factual statements and uncertainty expressions. In addition, we include #Incor to explicitly measure the absolute number of errors that the model is making.

For Uncertain Accuracy, we approximate $C_{\text{unc}}^{\text{true}}$ based on the following assumption: we regard a claim as truly uncertain if the model cannot directly answer it correctly in short-form QA. Similar to previous work (Tian et al., 2024; Zhang et al., 2024d; Farquhar et al., 2024), we first rewrite atomic facts into specific short-form questions. For example, the fact "[a company] was founded in 1996" is converted into the question "When was [a company] founded?". We present the short-form questions to the model and collect the model’s answer. If the model fails to answer a question correctly, it indicates that the model have high uncertainty about that fact and the claim can be regarded as truly uncertain. We use the number of questions that lead to wrong answers as the approximation of $C_{\text{unc}}^{\text{true}}$. Moreover, we focus on uncertainty about specific and granular details and exclude general and vague claims (as judged by GPT-4o; more discussion in Limitation Section).

In addition to FA and UA, we further include a metric of #Incor., which tracks the total number of incorrect statements generated by the model. While accuracy focuses on the ratio of correct to incorrect facts, the #Incor. metric directly measures the ab-

solute number of incorrect statements. Reducing such incorrect facts is an essential goal for improving the factuality of long-form responses, as a high number of incorrect statements directly undermines the overall reliability of the model’s output.

4.3 Baselines

In our experiments, we compare our LOGU-SFT and LOGU-DPO³ against four baseline methods (see implementation details of baselines in Appendix C.1): 1) **Unc-Zero**: The model is directly prompted to express uncertainty in its output if it is not sure about any claims. 2) **Unc-Few**: Building on Unc-Zero, we provide the model with an additional set of 10 hand-crafted QA examples where uncertainty is explicitly expressed in the answers, following an in-context learning approach. For each example, we provide the question paired with a long-form answer that includes uncertainty expressions. Each example is formatted as $\langle Q, A_{\text{unc}} \rangle$. 3) **Pair-Few**: Extending Unc-Few, we provide the model with both a long-form answer with only certain expressions A_{cert} and another one with uncertainty expressions A_{unc} for each query. Each example is formatted as $\langle Q, A_{\text{cert}}, A_{\text{unc}} \rangle$. The goal of including both A_{cert} and A_{unc} is to teach the model to differentiate between these two through in-context learning. 4) **Self-Refine** (Madaan et al., 2023): Instead of generating the response in one pass, we apply a draft-and-refine setup. The model is asked to first generate an initial response and then refine the uncertain facts into explicit uncertainty expressions in a second pass. We also compare our method with several post-hoc methods, where we apply a two-stage pipeline to first estimate the uncertainty for each claim and then revise the most uncertain claims (see Appendix C.2 and C.3).

5 Result and Analysis

5.1 Main Results

LoGU greatly enhances generation performance with higher accuracy and fewer incorrect statements. Table 1 presents our main experimental results. With both models, LoGU-DPO consistently achieves the best performance across all three datasets. Specifically, it improves accuracy from 38.8% to 65.4% for Mistral and from 51.9%

³In our experimental setting, all references to LOGU-DPO denote a two-stage training approach combining DPO and SFT. We observed that a model trained solely by DPO does not effectively express uncertainty, which significantly impairs its performance (as shown in Table 8).

Method	Bios			LongFact			WildHallu			
	FA↑	UA↑	#Incor↓	FA↑	UA↑	#Incor↓	FA↑	UA↑	#Incor↓	
Mistral-7B-Instruct	Orig.	38.8	–	27.9	86.2	–	4.55	71.5	–	8.31
	<i>Prompt-Based</i>									
	Unc-Zero	43.6	76.6	22.3	88.5	29.9	1.85	75.4	48.6	5.20
	Unc-Few	41.7	74.8	22.8	85.6	32.4	5.58	72.6	57.3	10.1
	Pair-Few	40.3	66.7	25.9	86.1	35.1	5.08	73.4	51.6	9.80
	Self-Refine	38.3	57.4	27.2	86.2	42.4	4.61	72.4	50.6	8.29
	<i>Training-Based</i>									
	LoGU-SFT	54.5	77.1	11.4	88.6	43.5	3.21	79.2	51.1	5.73
LoGU-DPO	65.4	80.7	6.54	91.3	54.6	<u>2.09</u>	84.4	61.8	3.49	
Llama3-8B-Instruct	Orig.	51.9	–	20.4	85.5	–	7.45	74.4	–	6.24
	<i>Prompt-Based</i>									
	Unc-Zero	53.8	65.4	14.9	89.3	30.4	2.67	78.7	45.8	3.29
	Unc-Few	58.7	69.1	11.7	88.3	40.6	4.01	80.2	48.6	3.44
	Pair-Few	60.6	46.3	9.86	86.7	42.2	4.78	84.0	48.1	2.33
	Self-Refine	51.8	40.5	18.8	85.0	27.2	6.38	77.0	28.7	5.13
	<i>Training-Based</i>									
	LoGU-SFT	58.5	69.7	11.5	87.5	44.8	4.53	78.9	44.6	5.37
LoGU-DPO	71.4	70.9	4.84	91.5	47.5	2.36	86.0	52.8	<u>2.68</u>	

Table 1: Comparison of the performance of LOGU-SFT and LOGU-DPO with baseline methods. The best results are **bolded**, and the second best ones are underlined.

to 71.4% for Llama3 on the *Bios* dataset. This demonstrates that our approach can effectively reduce hallucination and improve overall factuality.

LOGU-DPO enables the model to accurately express uncertainty. In all baseline methods, although the model can express some uncertainty in long-form generation, the uncertain accuracy remains low. This indicates that the models are merely mimicking uncertainty expressions without reflecting actual uncertainty. While LOGU-SFT encourages more uncertainty expressions, it does not guarantee that the expressed uncertainty is accurate. LOGU-DPO improves this uncertainty misalignment by training the model to express uncertainty only when it may genuinely make mistakes on a given fact, thus achieving better uncertain accuracy. Figure 5 provides a detailed case study of how LOGU-DPO improves over LOGU-SFT and the original model in generating more accurate uncertainty expressions.

LOGU-DPO shows generalization ability on Out-of-Domain Dataset. To assess the generalizability of our approach, we conduct additional experiments using the ASQA dataset (Stelmakh et al., 2022), which focuses on ambiguous factoid questions in long-form question answering. Since ASQA does not overlap with the three datasets (*Bios*, *LongFact*, *WildHallu*) used during training, it serves as an ideal out-of-domain benchmark for

Method	Mistral-7B			Llama3-8B		
	FA↑	UA↑	#Incor↑	FA↑	UA↑	#Incor↑
Orig.	71.3	–	4.00	77.9	–	3.88
Unc-Zero	73.4	60.7	1.73	79.3	51.7	1.73
Self-Refine	70.3	35.2	4.37	79.2	38.7	3.31
LoGU-SFT	79.9	63.2	2.49	82.4	48.8	3.17
LoGU-DPO	84.3	68.9	1.56	87.3	55.2	<u>2.02</u>

Table 2: Results on Out-of-Domain Dataset (ASQA). LOGU-DPO shows robustness when tested on dataset that differ significantly from the training data.

evaluating our model’s robustness. As shown in Table 2, our approach outperforms existing methods in ASQA, achieving higher scores in both Fact Accuracy (FA) and Uncertainty Accuracy (UA), further validating its generalizability.

5.2 Analysis

Different Training Data Sources. We explore cases where the training data are collected using the original responses from different source models. In Table 3, we train the Llama3-8B-Instruct model using data constructed from Mistral-7B-Instruct. We have the following two findings: 1) For LOGU-SFT, data from another LLM will not always lead to worse performance. This may be because LOGU-SFT primarily addresses the *uncertainty suppression* issue. As long as the training samples include uncertainty expressions, the model can

Method	FA	UA	#Incor
<i>Bios</i>			
LOGU-SFT	58.5	69.7	11.5
w/ mismatch data	46.8	69.6	14.4
LOGU-DPO	71.4	70.9	4.84
w/ mismatch data	69.2	68.1	5.88
<i>LongFact</i>			
LOGU-SFT	87.4	44.8	4.54
w/ mismatch data	86.7	43.2	3.94
LOGU-DPO	91.5	47.5	2.36
w/ mismatch data	89.9	44.1	3.16
<i>WildHallu</i>			
LOGU-SFT	78.9	44.6	5.36
w/ mismatch data	76.5	48.2	6.30
LOGU-DPO	86.0	52.8	2.68
w/ mismatch data	82.7	50.6	3.39

Table 3: Results of training Llama3-8B-Instruct with mismatch data (from Mistral-7B-Instruct) show that using the model’s own generations in the DPO improves its ability to express uncertainty properly.

learn the pattern. 2) For LOGU-DPO, using data from another LLM consistently leads to worse performance. We argue that this is because the DPO stage primarily serves to align uncertainty expression with unknown facts, which requires training data that reflects the intrinsic uncertainty of the specific model. Uncertainty for one model may not correspond to uncertainty for another. This finding underscores the importance of using a model’s own generations during the DPO stage.

Generation Statistics. Table 4 compares statistics of the original model’s generation against the LOGU-SFT and LOGU-DPO outputs. We observe that in *LongFact* and *WildHallu*, the total generation length and the total number of claims remain relatively consistent. The consistency demonstrates that our training strategies do not substantially reduce the amount of generated information. For *Bios*, the average number of tokens decreases considerably, from 400.6 to 255.9. This may be because the original *Bios* responses contain a high number of incorrect claims (27.9 on average), which are reduced after training. With LOGU-DPO, the number of incorrect claims drops significantly to 6.5, leading to much more accurate uncertainty expression. Overall, Table 4 confirms that our uncertainty training approach effectively preserves the total information content while significantly reducing the number of incorrect claims.

Uncertainty Categories. To gain a deeper understanding of what the model expresses as uncertainty, we conduct an in-depth analysis of uncer-

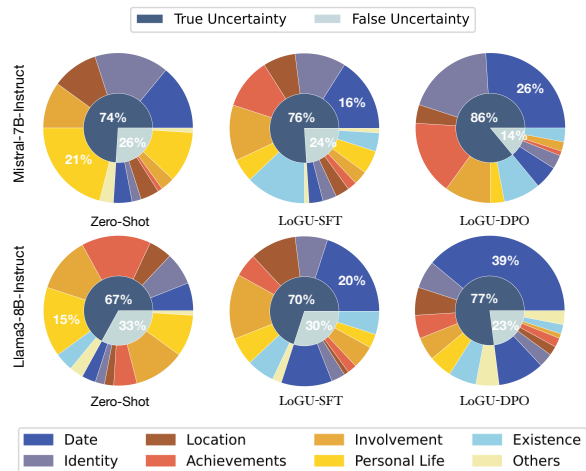


Figure 3: The distribution of uncertainty categories expressed by Zero-Shot, LOGU-DPO, LOGU-SFT on the *Bios* dataset. The inner layer represents the proportion of atomic claims classified as True Uncertainty (C_{unc}^{true}), while the outer layer shows the distribution across the eight uncertainty categories.

tainty categories with the *Bios* dataset, as shown in Figure 3. We first manually categorize the topics of uncertain expressions into eight categories that cover all atomic claims, and then use GPT-4o to label each uncertainty expression (details of the categories are provided in Appendix G). As illustrated in Figure 3, prompt-based methods tend to express true uncertainty in a broad manner. For instance, a significant proportion of uncertain expressions in the personal life category are similar to “I have no detailed information about Annabel’s early life.”. In contrast, training-based methods, particularly LOGU-DPO, express uncertainty in a more specific manner, often tied to concrete details, such as “I am not sure when Annabel was born.”. Date-related uncertainty constitutes 26% in Mistral and 39% in Llama3, while personal life uncertainty is minimal. This specificity in expressing uncertainty makes training-based methods more effective in reducing incorrect information about critical details in biographies, particularly dates and timelines, when relevant knowledge is missing. However, the general uncertainty expressed by prompt-based methods is less effective in preventing errors in key details, as it often refers to vague pieces of information. Regarding false uncertainty expressions, both approaches show balanced error distributions across categories, with no category dominating.

Method	#Cor	#Incor	#Unc	#Token
<i>Bios</i>				
Orig.	17.7	27.9	–	400.6
LoGU-SFT	13.6	11.4	3.82	276.8
LoGU-DPO	12.3	6.54	6.32	255.9
<i>LongFact</i>				
Orig.	28.4	4.55	–	282.6
LoGU-SFT	24.9	3.21	3.28	283.4
LoGU-DPO	21.8	2.09	5.35	277.9
<i>WildHallu</i>				
Orig.	20.8	8.31	–	211.9
LoGU-SFT	21.9	5.74	2.43	268.3
LoGU-DPO	18.9	3.49	4.50	266.7

Table 4: Statistics for the original (Orig.), LoGU-SFT, and LoGU-DPO responses using Mistral-7B-Instruct. #Cor, #Incor and #Unc indicate the average number of correct, incorrect and uncertain claims, respectively. #Token indicates the average token count per response.




Figure 4: Human evaluation results on helpfulness and fluency of responses generated using different methods.

5.3 Human Evaluation

We conduct a human evaluation to further evaluate the user experience on the uncertainty expression, considering two key aspects: helpfulness and fluency. *Helpfulness* measures the extent to which the model’s output aids the user in understanding the topic and making informed decisions. *Fluency* assesses the cohesion and coherency of the model’s output (see detailed information of the human evaluation are provided in Appendix H). We compare the original, Zero-Shot, and LoGU-DPO (with uncertainty ratio α values of 0.2 and 0.6) responses. As shown in Figure 4, LoGU-DPO with a 20% uncertainty ratio provides the most helpful responses. More importantly, annotators reported that using appropriate expressions of uncertainty can highlight information essential to the answers but the model does not know (such as a person’s birthdate in a biography), thereby complementing accurate statements. Although the fluency of the LoGU-DPO outputs was slightly lower than that of the original responses, it remained smooth overall. Additionally, the variance in both helpfulness and fluency scores for LoGU-DPO with a 20% uncer-

tainty ratio was smaller, indicating that our method produces more stable responses and provides users with a consistently positive experience. Moreover, increasing the uncertainty ratio to 60% results in a noticeable decline in both fluency and helpfulness, suggesting that a higher uncertainty ratio is not always better. More experiments on uncertainty ratio can be found in Appendix E.

6 Related Work

Fine-tuning LLMs for Factuality. There has been substantial work on fine-tuning LLMs to improve factuality. Tian et al. (2024) and Zhang et al. (2024d) proposed different approaches for annotating responses to fine-tune models using Direct Preference Optimization (DPO). The former focuses on leveraging external knowledge sources, while the latter emphasizes the utilization of the model’s own signals to rank the truthfulness of responses. Furthermore, Lin et al. (2024) underscores the importance of factual data in alignment training, employing distinct reward methods for factual and non-factual data to refine model performance. Moreover, multiple studies have highlighted that introducing new knowledge during fine-tuning can sometimes lead to increased hallucination (Lin et al., 2024; Gekhman et al., 2024; Kang et al., 2024). This insight suggests it may be more effective to let models express uncertainty rather than forcing them to state unfamiliar facts.

Teaching LLMs to Express Uncertainty. There are two main approaches to teaching LLMs to express uncertainty: prompt-based and training-based. Prompt-based methods (Tian et al., 2023; Xiong et al., 2024a) design specific prompts to encourage LLMs to convey uncertainty in their responses. On the other hand, training-based approaches (Lin et al., 2022; Li et al., 2024; Zhang et al., 2024c; Xu et al., 2024) develop specialized datasets to tune the models to express their uncertainty explicitly (e.g., "I am sure/unsure"). However, these methods focus on short-form question-answering, where each response contains limited facts. They also provide only one single estimation of uncertainty for the full response. Our work aims to address this gap by enabling models to selectively express uncertainty in responses involving multiple facts for long-form generation.

Long-form Uncertainty. In contrast to expressing uncertainty during generation, several studies focus on externally estimating model uncertainty.

For example, Zhang et al. (2024b) introduced LUQ, a consistency-based uncertainty estimation method for long-form generation. Huang et al. (2024) explored the calibration of long-form responses, correlating overall response quality with the associated confidence level. However, we argue that post-hoc uncertainty estimation incurs significantly higher computational costs during inference. Therefore, directly expressing uncertainty within the response offers a more efficient approach. To this end, Band et al. (2024) introduced the concept of linguistic calibration for long-form text generation, which aims to give a calibrated verbalized confidence estimation for each statement.

7 Conclusion

In this paper, we introduce the LoGU task, which aims to enhance the ability of LLMs to express uncertainty in long-form generation. To tackle the problems of Uncertainty Suppression and Uncertainty Misalignment, we propose a decomposition-based data collection framework and a two-stage training pipeline with LOGU-SFT and LOGU-DPO. Our experiments on three long-form QA datasets demonstrate that our approach significantly improves factual accuracy and uncertain precision, effectively reducing hallucinations without sacrificing the comprehensiveness of the generated content. These results underscore the importance of accurate uncertainty modeling in long-form generation, paving the way for more reliable and trustworthy LLM applications in real-world scenarios.

Limitation

This work focuses on enabling models to express uncertainty in factual QA datasets. Future research could explore uncertainty expression in other long-form generation tasks, such as grammar checking, summarization, and translation. Regarding factuality metrics, we do not report commonly used metrics for long-form QA tasks, such as ROUGE-L and recall, because the datasets used in our study (*Bios*, *LongFact*, *WildHallu*) do not provide gold-standard answers.

Furthermore, the uncertainty expression in our work is embedded in the text and treated as binary. We do not require models to indicate degrees of uncertainty (e.g., "I am 10% unsure about..."). We argue that binary expression is more human-interpretable and aligns better with natural communication. However, future work could explore

machine-interpretable formats, such as attaching confidence tags to each sentence.

In Section 4.2, we evaluate uncertainty expression on specific and granular details. Two under-explored areas remain: (1) how to probe known versus unknown knowledge: our current approach uses short-form QA, but more probing methods could be applied; and (2) how to define the specific and granular aspects to focus on; we currently rely on prompting GPT-4o for filtering, but other strategies could be investigated. Our follow-up work (Yang et al., 2025) further explores these two limitations in detail.

Ethics Statement

Our research adheres to rigorous ethical guidelines. We verified the licenses of all softwares and datasets we used in this study and ensured full compliance with their terms. During the human annotation process, all annotators provided informed consent for their data to be used as part of the project. No privacy concerns have been identified. Furthermore, we have thoroughly assessed the project and do not anticipate any additional potential risks.

Acknowledgement

This work is supported by the Chinese NSF Major Research Plan (No.92270121).

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. *Gpt-4 technical report*. *ArXiv preprint*, abs/2303.08774.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hongchang Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, et al. 2023. *Longbench: A bilingual, multitask benchmark for long context understanding*. *ArXiv preprint*, abs/2308.14508.
- Neil Band, Xuechen Li, Tengyu Ma, and Tatsunori Hashimoto. 2024. *Linguistic calibration of long-form generations*. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, et al. 2024. A survey on evaluation of large language models. *ACM Transactions on Intelligent Systems and Technology*, 15(3):1–45.

- 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](#). In *First Conference on Language Modeling*.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2021. [Editing factual knowledge in language models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6491–6506, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Ekaterina Fadeeva, Roman Vashurin, Akim Tsvigun, Artem Vazhentsev, Sergey Petrakov, 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.
- Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. 2024. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017):625–630.
- Zorik Gekhman, Gal Yona, Roei Aharoni, Matan Eyal, Amir Feder, Roi Reichart, and Jonathan Herzig. 2024. [Does fine-tuning llms on new knowledge encourage hallucinations?](#) *ArXiv preprint*, abs/2405.05904.
- Muhammad Usman Hadi, Rizwan Qureshi, Abbas Shah, Muhammad Irfan, Anas Zafar, Muhammad Bilal Shaikh, Naveed Akhtar, Jia Wu, Seyedali Mirjalili, et al. 2023. A survey on large language models: Applications, challenges, limitations, and practical usage. *Authorea Preprints*.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Yukun Huang, Yixin Liu, Raghuv eer Thirukovalluru, Arman Cohan, and Bhuwan Dhingra. 2024. [Calibrating long-form generations from large language models](#). *Preprint*, arXiv:2402.06544.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7B](#). *Preprint*, arXiv:2310.06825.
- Lekang Jiang, Caiqi Zhang, Pascal A. Scherz, and Stefan Goetz. 2025. [Can large language models generate high-quality patent claims?](#) In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 1272–1287, Albuquerque, New Mexico. Association for Computational Linguistics.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, Deep Ganguli, Danny Hernandez, Josh Jacobson, Jackson Kernion, Shauna Kravec, Liane Lovitt, Kamal Ndousse, Catherine Olsson, Sam Ringer, Dario Amodei, Tom Brown, Jack Clark, Nicholas Joseph, Ben Mann, Sam McCandlish, Chris Olah, and Jared Kaplan. 2022. [Language models \(mostly\) know what they know](#).
- Katie Kang, Eric Wallace, Claire Tomlin, Aviral Kumar, and Sergey Levine. 2024. [Unfamiliar finetuning examples control how language models hallucinate](#). *ArXiv preprint*, abs/2403.05612.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- 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.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pages 611–626.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- Jiaqi Li, Yixuan Tang, and Yi Yang. 2024. [Know the unknown: An uncertainty-sensitive method for llm instruction tuning](#).
- Sheng-Chieh Lin, Luyu Gao, Barlas Oguz, Wenhan Xiong, Jimmy Lin, Scott Yih, and Xilun Chen. 2024. [FLAME: Factuality-aware alignment for large language models](#). In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [Teaching models to express their uncertainty in words](#). *ArXiv preprint*, abs/2205.14334.

- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhunoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. [Self-refine: Iterative refinement with self-feedback](#). 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*.
- Meta. 2024. [Llama 3 model card](#).
- 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.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). 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*.
- 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.
- Katherine Tian, Eric Mitchell, Huaxiu Yao, Christopher D. Manning, and Chelsea Finn. 2024. [Fine-tuning language models for factuality](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher Manning. 2023. [Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5433–5442, Singapore. Association for Computational Linguistics.
- Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Jie Huang, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024. [Long-form factuality in large language models](#). In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*.
- Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. 2024a. [Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. 2024b. [Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Tianyang Xu, Shujin Wu, Shizhe Diao, Xiaozhe Liu, Xingyao Wang, Yangyi Chen, and Jing Gao. 2024. [Sayself: Teaching llms to express confidence with self-reflective rationales](#). *Preprint*, arXiv:2405.20974.
- Ruihan Yang, Caiqi Zhang, Zhisong Zhang, Xinting Huang, Dong Yu, Nigel Collier, and Deqing Yang. 2025. [Uncle: Uncertainty expressions in long-form generation](#). *Preprint*, arXiv:2505.16922.
- Tao Yuan, Xuefei Ning, Dong Zhou, Zhijie Yang, Shiyao Li, Minghui Zhuang, Zheyue Tan, Zhuyu Yao, Dahua Lin, Boxun Li, et al. 2024. [Lv-eval: A balanced long-context benchmark with 5 length levels up to 256k](#). *ArXiv preprint*, abs/2402.05136.
- Caiqi Zhang, Zhijiang Guo, and Andreas Vlachos. 2024a. [Do we need language-specific fact-checking models? the case of Chinese](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1899–1914, Miami, Florida, USA. Association for Computational Linguistics.
- Caiqi Zhang, Fangyu Liu, Marco Basaldella, and Nigel Collier. 2024b. [Luq: Long-text uncertainty quantification for llms](#). *Preprint*, arXiv:2403.20279.
- Hanning Zhang, Shizhe Diao, Yong Lin, Yi Fung, Qing Lian, Xingyao Wang, Yangyi Chen, Heng Ji, and Tong Zhang. 2024c. [R-tuning: Instructing large language models to say ‘I don’t know’](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7113–7139, Mexico City, Mexico. Association for Computational Linguistics.
- Xiaoying Zhang, Baolin Peng, Ye Tian, Jingyan Zhou, Lifeng Jin, Linfeng Song, Haitao Mi, and Helen Meng. 2024d. [Self-alignment for factuality: Mitigating hallucinations in LLMs via self-evaluation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1965, Bangkok, Thailand. Association for Computational Linguistics.

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](#). *Preprint*, arXiv:2303.18223.

Wenting Zhao, Tanya Goyal, Yu Ying Chiu, Liwei Jiang, Benjamin Newman, Abhilasha Ravichander, Khyathi Chandu, Ronan Le Bras, Claire Cardie, Yuntian Deng, and Yejin Choi. 2024. [Wildhallucinations: Evaluating long-form factuality in llms with real-world entity queries](#).

Appendix

A Instruction Prompt Examples

The instruction prompts for the revision and assembling procedures in §3 are presented in Listing 1.

Listing 1: The instruction prompts for Revision and Assembling procedures.

```
Revision Instruction:
You will be given a series of atomic facts, each labeled as ##certain## or ##
uncertain##. Please adjust each fact following these steps:
- For facts labeled as ##certain##, leave them unchanged.
- For facts labeled as ##uncertain##, adjust them to express uncertainty without
focusing on overly specific details. Instead of being uncertain about exact facts,
use more general phrases like 'I am not sure when/where/how/what' to convey the
uncertainty.

Here are some uncertainty expressions you can refer to:
- It is uncertain/unclear/not sure/not known
- I am uncertain/unclear/not sure/not known
- There is no information about
- There is no concrete answer about
...

Output each fact(including unchanged facts labeled #ceratin# and adjusted facts with
#uncertain# label) in order, as a single line beginning with ###.

For Example:

Facts:
Kang Ji-hwan was born on March 16, 1982. ##uncertain##
Kang Ji-hwan was born in Seoul, South Korea. ##uncertain##
Chief Jones is a respected figure. ##uncertain##
Chief Jones is a respected figure in law enforcement. ##uncertain##
He has had a successful career in the police force. ##uncertain##
He rose through the ranks. ##certain##

Output:
### I do not know when Kang Ji-hwan was born.
### I am not sure where Kang Ji-hwan was born.
### It is unclear whether Chief Jones is a respected figure.
### It is uncertain whether Chief Jones is a respected figure in law enforcement.
### It is uncertain whether he might had a successful career in the police force.
### He rose through the ranks.

Now it's your turn to answer:

Assembling Instruction:
Your task is to concatenate a provided list of atomic facts, each articulated with
either certainty or uncertainty, into a cohesive narrative following the guidelines
below:

- All facts in the list, regardless of their certainty, MUST BE included in the
generated text. Eliminate any duplicates that may exist.
- Refrain from adding any facts that are not mentioned in the original atomic facts
list.
- Your narrative must flow smoothly and coherently. Uncertainty (if present in the
atomic facts list) should be expressed diversely, employing phrases such as 'might',
'it is/I am not sure'.

Here is the atomic facts:

{paragraph}

Output:
```

B Implementation Details.

In our experiments, we utilize the Llama3-8B-Instruct (Meta, 2024) and Mistral-7B-Instruct (Jiang et al., 2023) as our base models.

Training Data The LOGU-SFT and LOGU-DPO datasets comprise a mix of Bios, WildHallu, and LongFact. Here, we used an uncertainty ratio of $\alpha = 0.2$ to construct our training data. To enrich the training data and enhance the model’s robustness, we implemented data augmentation on the instructions for Bios and WildHallu. We applied five specific transformations as follows:

```
Original Prompt:
Bios: Tell me a bio of
WildHallu: In a paragraph, can you tell me about
```

```
Data Augmentation:
1. Tell me what you know about
2. Can you provide a detailed introduction of
3. Can you tell me about
4. Can you provide information about
```

The specific data set size during the LOGU-SFT and LOGU-DPO phases for Mistral-7B-Instruct and Llama3-8B-Instruct is detailed in Table 5. We also provide the examples of our LOGU-SFT and LOGU-DPO training data as below:

LoGU-SFT Example:

```
{
  "instruction": "Tell me a bio of Suthida.",
  "input": "",
  "output": "Suthida Vajiralprabha is the current queen consort of Thailand. Before her ascension to this prestigious role, she was born into a middle-class family and rose through the ranks in the Royal Thai Police, where she was known for her dedication and hard work. Although it is uncertain, some say she might have worked as a nurse, but I am not sure where she might have done so in Bangkok.\n\nIn a surprising turn of events, Suthida Vajiralprabha married King Maha Vajiralongkorn of Thailand on May 1, 2019. Following her marriage, she was given the title of Queen Consort and has since taken on various royal duties, often seen attending official events with the king and engaging in charitable work.\n\nDespite her royal status, Suthida Vajiralprabha is known for her simple and down-to-earth personality. Her humility and dedication to serving the people have won the hearts of many Thais. While I do not know what her specific royal title is, I can affirm that she is fluent in Thai and continues to be a beloved figure in Thailand."
},
```

LoGU-DPO Example:

```
{
  "instruction": "Who is Gottfried Wilhelm Leibniz?",
  "input": "",
  "chosen": "Gottfried Wilhelm Leibniz, born on July 1, 1646, and passing away on November 14, 1716, was a German philosopher, mathematician, and polymath. Widely regarded as one of the greatest thinkers of the late 17th century, Leibniz made significant contributions to various fields including metaphysics, logic, theology, and particularly mathematics. Notably, he developed calculus independently of Sir Isaac Newton and invented the Leibniz wheel, although I am not sure whether he also invented the Leibniz calculus.\n\nLeibniz is perhaps most famous for his philosophical concept of monadology, which posits that the universe is made up of individual, simple substances called monads. While he was undeniably a prolific writer, the exact number of letters and manuscripts he produced remains uncertain. \n\nAs a member of several European academies, Leibniz held various positions of influence throughout Europe. However, it is not clear whether he made significant contributions to the fields of physics and diplomacy. Overall, Leibniz's legacy as a thinker who traversed multiple disciplines has ensured his lasting impact on the intellectual landscape.",
  "rejected": "Gottfried Wilhelm Leibniz (July 1, 1646 \u2013 November 14, 1716) was a German philosopher, mathematician, and polymath who is considered one of the greatest thinkers of the late 17th century. He is best known for his contributions to philosophy, particularly in the areas of metaphysics, logic, and theology. Leibniz is also known for his work in mathematics, including his development of the calculus independent of Sir Isaac Newton, and his invention of the Leibniz wheel and the Leibniz calculus.\n\nLeibniz is perhaps most famous
```

for his philosophical concept of monadology, which posits that the universe is made up of individual, simple substances called monads. He also made significant contributions to the fields of physics, law, and diplomacy. Leibniz was a prolific writer and is believed to have produced over 4,000 letters and manuscripts during his lifetime. He was a member of several European academies and held various positions of power and influence throughout Europe."

Stage	<i>bios</i>	<i>LongFact</i>	<i>WildHallu</i>	Total
<i>Mistral-7B-Instruct</i>				
LOGU-SFT	7970	29160	1548	38678
LOGU-DPO	3555	19096	911	23562
<i>Llama3-8B-Instruct</i>				
LOGU-SFT	8475	27365	1417	37257
LOGU-DPO	3933	19096	851	23880

Table 5: Training dataset size for LOGU-SFT and LOGU-DPO.

Fine-tuning Details We run LOGU-DPO and LOGU-SFT experiments with 8 NVIDIA A100-40GB GPUs. We conduct experiments with the LlamaFactory code base⁴. During the LOGU-SFT phase, we fine-tune the base models on a mixed dataset of 40k. In the LOGU-DPO phase, we continue fine-tuning the models on a mixed dataset of 20k instances. Building upon prior research, which highlights the MLP layer as a crucial element for embedding knowledge within the LLM transformer architecture (De Cao et al., 2021), we only fine-tune the weight matrix of the attention layer using LoRA (Hu et al., 2022). This method allows us to adjust the model’s ability to express knowledge boundaries without altering its internal knowledge structure. The configurations of our hyper-parameters are detailed in Table 6.

Configuration	LOGU-SFT	LOGU-DPO
Model	Mistral-7B(Llama3-8B)-Instruct	Mistral-7B(Llama3-8B)-Instruct
Number of epochs	3	3
Total Batch size	64 samples	64 samples
Optimizer	Adam (Kingma and Ba, 2015) ($\beta_1 = 0.9, \beta_2 = 0.98, \epsilon = 1 \times 10^{-8}$)	Adam (Kingma and Ba, 2015) ($\beta_1 = 0.9, \beta_2 = 0.98, \epsilon = 1 \times 10^{-8}$)
Learning rate	5×10^{-5}	1×10^{-5}
Warmup Ratio	0.1	0.1
LoRA Target	q_{proj}, v_{proj}	q_{proj}, v_{proj}
LoRA Parameters	$r = 8, \alpha = 16, dropout = 0.05$	$r = 8, \alpha = 16, dropout = 0.05$
Training Time	1h 31m 30s (1h 43m 3s)	1h 55m (1h 3m 8s)

Table 6: Fine-tuning hyper-parameters for LOGU-SFT and LOGU-DPO.

Evaluation We use vLLM (Kwon et al., 2023) for our LLM inference tasks, with the following parameters: temperature = 0.7, top- p = 0.95, and a maximum output of 1024 tokens. For Fact-Checking, Uncertainty Revision, Assemble procedures, we set the temperature to 0. We use GPT-4o as the auxiliary model for generating atomic claims and fact-checking the LLM.

⁴<https://github.com/hiyouga/LLaMA-Factory>

C Baselines

C.1 Generative Methods

Here we list the prompts of our baselines (zero, few-shot, paired few-shot and self-refine) in §4.3.

Listing 2: The instruction prompts of key procedures.

Zero Shot(Bios):

Your task is to write a biography for a specific entity. You should express uncertainty for any information you are not familiar with.

Zero Shot(LongFact):

Your task is to answer the given question about a specific object (e.g., a person, place, event, company, etc.). Express uncertainty for any information you are not familiar with (e.g., 'I am not sure if/whether', 'It is uncertain that').

Zero Shot(WildHallu):

Your task is to write a paragraph for a specific entity. You should express uncertainty for any information you are not familiar with (e.g., 'I am not sure if/whether', 'It is uncertain that').

Few Shot Examples:

Question: In a paragraph, could you tell me what you know about Vycta.co.ke?

Answer: Vycta.co.ke is a Kenyan-based digital health platform that provides healthcare services to patients. Additionally, Vycta.co.ke offers various health-related services, ensuring comprehensive support for its users.

Paired Few Shot Examples:

Question: In a paragraph, could you tell me what you know about Vycta.co.ke?

Good Answer: Vycta.co.ke is a Kenyan-based digital health platform that provides healthcare services to patients. Additionally, Vycta.co.ke offers various health-related services, ensuring comprehensive support for its users.

Bad Answer: Vycta.co.ke is a Kenyan-based digital health platform that provides a range of healthcare services to patients and healthcare providers. The platform offers a telemedicine service that allows patients to consult with doctors remotely, as well as a health insurance platform that provides access to affordable healthcare coverage. Vycta also offers a range of health-related products and services, including a health monitoring app that allows patients to track their health metrics and receive personalized health advice. The platform aims to increase access to healthcare services, particularly in rural and underserved areas, and to make healthcare more affordable and convenient for patients.

Self Refine:

Revise your previous response to the question by adjusting the expression of certainty.

For any information you are unsure about, frame your response to reflect this uncertainty. Keep the expression of information you are certain about unchanged.

For example:

Question: Tell me a bio of Miguel Angel Felix Gallardo.

Answer: Miguel Angel Felix Gallardo, also known as "El Padrino" (The Godfather), is a Mexican drug lord and former leader of the Guadalajara cartel, one of the most powerful drug trafficking organizations in Mexico during the 1980s. He was born on August 7, 1946, in La Barranca de Agua Caliente, Sinaloa, Mexico. Gallardo began his criminal career in the 1960s, working as a small-time smuggler and drug dealer. He quickly rose through the ranks, becoming a key player in the Mexican drug trade and eventually forming his own organization, the Guadalajara cartel, in the 1970s.

Refine: Miguel Angel Felix Gallardo, who is widely known as 'El Padrino' (The Godfather), is a notorious Mexican drug lord. He emerged as a prominent figure in the criminal underworld by forming the Guadalajara cartel in the 1970s. Gallardo's rise to power was marked by his adept maneuvering within the Mexican drug trade, where he quickly became a key player.

Now it is your turn to refine:

Question: {question}

Answer: {answer}

Refine:

C.2 Post-hoc Methods

We incorporate two baselines—P(true) (Kadavath et al., 2022) and Semantic Entropy (SE) (Kuhn et al., 2023)—which rely on post-hoc uncertainty estimates derived from the model’s internal signals. For each baseline, we compute a post-hoc uncertainty score (P(true) or SE) for each atomic claim. Based on these scores, we rank the atomic claims within each dataset and set a threshold for uncertain claims (we chose the bottom 10%). These claims were then modified using GPT-4 to explicitly express uncertainty.

C.3 Comparison of LOGU-SFT, LOGU-DPO and Post-hoc Methods.

In this section, we compare the post-hoc baselines discussed in §C.2 with our LOGU-SFT and LOGU-DPO methods. Similar to the main experiments, we evaluate on three datasets: *Bios*, *LongFact*, and *WildHallu*, using Mistral-7B as the backbone model. The evaluation metrics are Factual Accuracy (FA) and Uncertainty Accuracy (UA). As shown in Table 7, our LoGU-DPO outperforms the post-hoc uncertainty baselines (P(true) and SE) in both FA and UA across all three datasets. We hypothesize that this ineffectiveness stems from the accuracy limitations of these uncertainty estimation methods in the context of long-text generation.

Method	Bios		LongFact		WildHallu	
	FA↑	UA↑	FA↑	UA↑	FA↑	UA↑
Orig.	38.8	–	86.2	–	71.5	–
P(true)	45.1	72.6	87.4	26.8	73.8	44.2
SE	41.7	83.1	88.4	36.3	73.4	53.1
LOGU-SFT	54.5	77.1	88.6	43.5	79.2	51.1
LoGU-DPO	65.4	80.7	91.3	54.6	84.4	61.8

Table 7: Comparison of LOGU-SFT, LOGU-DPO and Post-hoc Methods.

D Comparison of LOGU-DPO and DPO-Only.

In this section, we compare the performance of LOGU-DPO and DPO-Only. Similar to the main experiments, we evaluate the models on three datasets: *Bios*, *LongFact*, and *WildHallu*, using Factual Accuracy (FA) and Uncertainty Accuracy (UA) as metrics. As shown in Table 8, we find that LOGU-DPO achieves the best performance, while training with DPO alone significantly decreases model performance, with UA metrics even falling behind LOGU-SFT.

	Method	Bios		LongFact		WildHallu	
		FA↑	UA↑	FA↑	UA↑	FA↑	UA↑
Mistral-7B	Orig.	38.8	–	86.2	–	71.5	–
	LoGU-SFT	54.5	77.1	88.6	43.5	79.2	51.1
	LoGU-DPO	65.4	80.7	91.3	54.6	84.4	61.8
	DPO-only	68.7	76.6	91.8	46.8	83.4	52.2
Llama3-8B	Orig.	51.9	–	85.5	–	74.4	–
	LoGU-SFT	58.5	69.7	87.5	44.8	78.9	44.6
	LoGU-DPO	71.4	70.9	91.5	47.5	86.0	52.8
	DPO-only	74.8	62.4	90.1	42.1	83.6	46.7

Table 8: Comparison of LOGU-DPO and DPO-Only.

E Impact of the Uncertainty Ratio Parameter α .

To investigate the impact of the uncertainty ratio (α) in the training data on the final results, we train the Mistral-7B-Instruct using training data with different values of α during the SFT stage. As shown in Table 9, the results indicate that $\alpha = 0.2$ yields the highest Fact Accuracy (FA) and Uncertainty Accuracy (UA) across the three datasets, highlighting a trade-off between uncertainty expression and response quality.

Method	Bios		LongFact		WildHallu	
	FA↑	UA↑	FA↑	UA↑	FA↑	UA↑
LOGU-SFT ($\alpha = 0.2$)	54.5	77.1	88.6	43.5	79.2	51.1
LOGU-SFT ($\alpha = 0.4$)	51.8	74.0	87.4	42.6	79.7	46.0
LOGU-SFT ($\alpha = 0.6$)	52.9	71.6	87.3	43.6	78.4	48.4
LOGU-SFT ($\alpha = 0.8$)	47.0	73.7	87.5	42.3	78.3	49.1

Table 9: Comparison of LOGU-SFT, LOGU-DPO and Post-hoc Methods.

F Case Study

We randomly selected a sample from the *WildHallu* dataset to compare Original Answer and LOGU-DPO Answer. Each atomic claim in the generated responses was manually fact-checked using Wikipedia and other sources accessible via Google. Correct statements were highlighted in green, incorrect statements in red, and uncertainty expressions in blue.

As shown in Figure 5, the original answer contains a mix of correct and incorrect statements, which diminishes the overall factual accuracy of the generated content. In the case of LOGU-SFT, it significantly reduces the number of incorrect statements. However, it also introduces uncertainty in the expression of some correct facts. For example, in the original answer, Aegon the Conqueror’s birth time is correctly stated as “27BC”, but in the LOGU-SFT answer, this fact is expressed with uncertainty. In contrast, the LOGU-DPO response preserves the correct statements while replacing incorrect ones with expressions of uncertainty. For instance, the original answer incorrectly identifies Aegon the Conqueror’s parents as “He was the elder son of Lord Jaxartes Targaryen of Dragonstone and Princess Rhaenyra Targaryen.” However, the LOGU-DPO response addresses this uncertainty by stating, “The identity of his parents remains uncertain.”, thereby significantly improving the overall accuracy of the response.

Prompt: In a paragraph, could you tell me what you know about Aegon the Conqueror?

Original Answer: Aegon Targaryen, also known as Aegon the Conqueror, was the first king of the Targaryen Dynasty in Westeros, **was born in 27 BC**, ruling from 278 AC to 300 AC in the history of George R.R. Martin's "A Song of Ice and Fire" series and the "Game of Thrones" books. He was the elder son of Lord Lord Jaxartes Targaryen of Dragonstone and Princess Rhaenyra Targaryen, the sister of King Jaehaerys I Targaryen. Aegon was born on Dragonstone and raised in exile after his grandmother, Queen Rhaenys Targaryen, was passed over for the throne in favor of her brother. When Aegon came of age, he led a military campaign to reclaim the Seven Kingdoms for House Targaryen, using his two dragons, Balerion, Vhagar. He landed in the Stormlands with his sister-wives and quickly gained allies through his victories and the fear of his dragons. [...] He took King's Landing and **was crowned King of the Andals and the First Men**, Lord of the Seven Kingdoms, and Protector of the Seven Kingdoms.

LoGU-SFT Answer: Aegon Targaryen, also known as Aegon the Conqueror, was a legendary figure in the history of the Seven Kingdoms, renowned for his conquests and the establishment of the Targaryen dynasty. Born in Valyria, a powerful city-state located in Essos, Aegon was the founder of the Targaryen dynasty and the first Targaryen king of Westeros. He is remembered as a conqueror who united the Seven Kingdoms under his rule, although **I am not sure about the exact time of his birth**. Aegon was a skilled warrior, capable of taming Balerion and using it to terrorize his enemies. [...] However, I am not sure how Aegon conquered the Riverlands or the North, and I also have uncertainties about how he defeated Harren the Black, the last Storm King. Nevertheless, his conquests culminated in the Battle of the Field of Fire, where he defeated the combined armies of the Reach and the Rock, solidifying his control over the realm.

LoGU-DPO Answer: Aegon Targaryen, renowned as Aegon the Conqueror, was the legendary first Targaryen king of Westeros. He is celebrated for leading the dragon-riding Targaryen conquering of the Seven Kingdoms, an event that marked the beginning of the Targaryen dynasty's long reign over the realm. Aegon **was born in Dragonstone in 27 BC**, though I am not sure when exactly this event occurred. The identity of his parents remains uncertain as well. Aegon's remarkable journey began when he launched an invasion of Westeros from Dragonstone with his sister-wife, Rhaenys Targaryen, and **two dragons**, Balerion, and Vhagar. This invasion was a pivotal moment in Westeros' history, as the Seven Kingdoms were at that time warring among themselves. [...] After his conquest, Aegon Targaryen was hailed as a great king and a powerful ruler. However, I am not sure when he died or how long he reigned.

Figure 5: Case Study: Green indicates factually correct statements, red represents incorrect ones, and blue highlights uncertain expressions. The Original Answers includes a mix of correct and incorrect statements, while the LOGU-DPO answers retains the correct statements and revises the incorrect ones into expressions of uncertainty.

G Uncertainty Categories

Detailed descriptions of eight categories derived from uncertainty in atomic claims.

- 0: **Date and Timing Uncertainty.**
This label applies to uncertainties about the specific times of events, important dates in a person's life, or any time-related facts.
Example: "When Marie Alexandrine Becker received the Lasker-DeBakey Clinical Medical Research Award is [/uncertain/]." or "The timing of when Bella Akhmadulina was awarded the Lenin Komsomol Prize for Literature is [/uncertain/]."
- 1: **Identity and Occupation Uncertainty.**
Use this for uncertainties about the real existence or the professional role of individuals.
Example: "It is [/uncertain/] whether Chief Jones is a real person." or "Whether Ravi is a fellow of the IACR is [/uncertain/]."
- 2: **Location and Geography Uncertainty.**
This covers uncertainties related to the places of birth, education, residence, or any geographical locations.
Example: "The place where Sara Paxton was born is [/uncertain/]." or "The place where Lousteau earned his undergraduate degree is [/uncertain/]."
- 3: **Achievements and Contributions Uncertainty.**
This involves uncertainties about someone's achievements, awards, or any professional contributions.
Example: "Whether Antonio Gasalla received the Premio Regione Piemonte is [/uncertain/]." or "The awards won by Wilfredo Vargas is [/uncertain/]."
- 4: **Involvement and Action Uncertainty.**
Pertains to uncertainties about someone's participation in projects, campaigns, or actions.
Example: "Liam Payne's involvement with the 'Text Santa' campaign is [/uncertain/]." or "It's unclear whether Mauro Icardi was loaned to Sampdoria."
- 5: **Personal Life and Background Information Uncertainty.**
Use this for uncertainties about personal details, family background, or other private aspects.
Example: "Kourosh Zolani's personal life is [/uncertain/]." or "Many aspects of Rivera's early life are [/uncertain/]."
- 6: **Existence and Factual Veracity Uncertainty.**
This label is for uncertainties about the factual existence of events, appearances in media, or historical facts.
Example: "It is [/uncertain/] whether Marianne McAndrew has appeared in 'Law & Order: Special Victims Unit: The Movie.'" or "It is uncertain whether Virginia Valli appeared in 'The Big Combo'."
- 7: **Others.**
Anything that does not fit into the categories above.

H Human Evaluation

We randomly selected 100 queries from the *bios*, *WildHallu*, and *LongFact* datasets. Human annotators were instructed to assess the usefulness and fluency of the responses generated by Mistral-7B using different methods. We defined these dimensions on a scale from 1 to 5. We recruited three students with bachelor’s degrees in computer science and fluent English to conduct the annotations. These annotators were not involved in our project and had no prior discussions related to it. To measure inter-annotator agreement, we used Fleiss’ Kappa, obtaining a score of 0.759, which indicates substantial agreement (approaching the “almost perfect” range of 0.8-1.0; Landis and Koch, 1977). The annotators were compensated above the local minimum hourly wage. The instruction to guide human annotators to evaluate the **Helpfulness** and **Fluency** of Long-Form Responses are present in the listing 3 and 4.

Listing 3: Instructions for Human Annotators: Evaluating the Helpfulness of Long-Form Responses

Your task is to evaluate the helpfulness of each long-form response based on the proportion of **Correct**, **Incorrect**, and **Uncertain** information. Each long-form response will be highlighted to indicate these categories for different atomic claims. Your goal is to assess how these proportions impact the user's ability to understand the topic and make informed decisions.

Definition Of Helpfulness:

Helpfulness refers to how effectively the response provides **useful**, **accurate**, and **relevant** information. A long-form response typically contains multiple atomic claims that can be categorized into:

- **Correct Information**: Accurate and relevant facts or details that directly address the query. The more correct information a response contains, the more useful and reliable it is for the user.
- **Incorrect Information**: Misleading, inaccurate, or irrelevant details that detract from the response's usefulness. **Incorrect** information can confuse the user or lead to misunderstandings, which reduces the overall helpfulness.
- **Uncertain Information**: Statements that express ambiguity, limited knowledge, or lack of confidence. Uncertainty can be **Helpful** if it highlights the limits of knowledge or provides context, and it is generally better than giving incorrect information for critical points. However, if it is vague, unexplained, or excessive, it may reduce the clarity and impact of the response.

Rating Scale:

You will rate the helpfulness of each long-form response on a scale from 1 to 5, depending on the balance of **Correct**, **Incorrect**, and **Uncertain Information**:

1: Not Helpful

- The response is dominated by incorrect information or excessive, vague uncertainty, making it mostly misleading or confusing. **Correct** information, if present, is minimal and not useful.

2: Minimally Helpful

- The response contains some correct information, but it is heavily outweighed by significant incorrect claims or unclear uncertainty. While there are elements of usefulness, the overall response may confuse or mislead the user.

3: Moderately Helpful

- The response presents a mixture of correct, uncertain, and incorrect information. The majority of the response is correct and relevant, but incorrect or unclear claims detract from its overall helpfulness.

4: Helpful

- The response is mostly correct, with only minor incorrect details or justified uncertainty. Uncertainty is presented clearly and contributes to the user's understanding. The correct information outweighs any minor errors or vagueness, making the response helpful overall.

5: Very Helpful

- The response is highly informative, with a high proportion of correct information and well-justified uncertainty where appropriate. **Incorrect** information is minimal

or absent. Even if uncertainty is present, it enhances understanding and complements the value of the correct information.

Listing 4: Instructions for Human Annotators: Evaluating the Fluency of Long-Form Responses

Your task is to evaluate the fluency of each long-form response. **Fluency** refers to the readability, flow, and naturalness of the language used in the response. You should assess the fluency based on overall how smoothly the response reads and whether it is free from awkward phrasing or unnatural language. You should not only judge the fluency purely on the presence of uncertain expressions.

Definition Of Fluency:

Fluency refers to how naturally and clearly the response is written. A fluent response should be well-structured, easy to read, and free of grammatical errors or awkward phrasing. **Fluency** is important in maintaining the reader's engagement and ensuring the response is easily understood.

Rating Scale:

You will rate the fluency of each long-form response on a scale from **1 to 5**, where the score reflects how smoothly and naturally the response reads.

1: Not Fluent

- The response is difficult to read due to frequent awkward phrasing, or disjointed structure. It requires significant effort to follow, and the meaning may be unclear.

2: Minimally Fluent

- The response is somewhat readable, but it contains frequent awkward phrasing that disrupt the flow. The meaning is mostly clear, but the language is clunky and requires effort to process.

3: Moderately Fluent

- The response is generally fluent, but there are some instances of awkward phrasing. The meaning is clear, but the overall readability could be improved.

4: Fluent

- The response is well-written, with only minor instances of slightly awkward phrasing. The response flows naturally, and the meaning is clear without any interruptions to the reader's understanding.

5: Very Fluent

- The response reads very smoothly, with no awkward phrasing. It is natural, clear, and well-structured, making it very easy for the reader to follow.