

DEBATEQA: Evaluating Question Answering on Debatable Knowledge

Rongwu Xu^{1,3*}, Xuan Qi^{1*},
Zehan Qi¹, Wei Xu¹, Zhijiang Guo²

¹ Tsinghua University, ² University of Cambridge, ³ Fangcun AI

Emails: {xrw22, qi-x22}@mails.tsinghua.edu.cn

* Equal contribution

Abstract

The rise of large language models (LLMs) has enabled us to seek answers to inherently debatable questions on LLM chatbots, necessitating a reliable way to evaluate their ability. However, traditional QA benchmarks assume fixed answers are inadequate for this purpose. To address this, we introduce **DEBATEQA**, a dataset of 2,941 debatable questions, each accompanied by multiple human-annotated partial answers that capture a variety of perspectives. We develop two metrics: *Perspective Diversity*, which evaluates the comprehensiveness of perspectives, and *Dispute Awareness*, which assesses if the LLM acknowledges the question’s debatable nature. Experiments demonstrate that both metrics are aligned with human preferences and stable across different underlying models. Using DEBATEQA with two metrics, we assess 12 prevalent LLMs and retrieval-augmented generation methods. Our findings reveal that while LLMs generally excel at recognizing debatable issues, their ability to provide comprehensive answers encompassing diverse perspectives varies considerably.



<https://github.com/pillsofwind/DebateQA>

1 Introduction

How often do you query a chatbot about a debatable issue? Questions like “Does Donald Trump have a terrible character?” or “How do crop circles form?” frequently arise in everyday life, reflecting human beings’ natural curiosity about topics that inherently lack fixed answers (Lowry and Johnson, 1981; Brady, 2009). With the advent of large language models (LLMs; OpenAI 2023b; Chowdhery et al. 2023; Touvron et al. 2023), we now turn to these models to seek “*proper*” answers to such questions. Evaluating the performance of language models on debatable queries is crucial for understanding and enhancing their capabilities.

Inherent difficulties hinder our way to assess model performance on such questions. Traditional question-answering (QA) benchmarks are typically designed to provide fixed answers to questions, as in datasets like SQuAD (Rajpurkar et al., 2016) and Natural Questions (Kwiatkowski et al., 2019). Even in scenarios where multiple legitimate answers are possible, such as TruthfulQA (Lin et al., 2022), the answers are often presented as multiple-choice questions, limiting the space of responses. Recent years have seen the emergence of long-form QA evaluations, such as ELI5 (Fan et al., 2019) and ASQA (Stelmakh et al., 2022a), which allow for more elaborate answers. However, these works have not focused on inherently debatable questions. The most related work to us is DELPHI (Sun et al., 2023), which curates a dataset with controversial questions, however, their emphasis is on identifying controversy rather than delving deeper into evaluating models’ responses.

What constitutes a proper answer when asked about debatable knowledge? Drawing inspiration from interdisciplinary literature, we propose two desirable properties for such answers. Firstly, the answer should inform the receiver that the issue at hand is debatable (Misco, 2011). Secondly, a proper answer should be comprehensive and include diverse perspectives, which are crucial for maintaining an atmosphere of neutrality, especially in public-related inquiries (Habermas, 1991). Motivated by these considerations, we introduce the **DEBATEQA** dataset, which comprises 2,941 manually annotated debatable questions. To rigorously evaluate the aforementioned properties of model answers, each question is paired with several *partial answers*, each reflecting a single viewpoint on the debatable issue. These partial answers are generated through a three-stage pipeline and then annotated by three annotators. Based on the partial answers in DEBATEQA, we propose the metric of **Perspective Diversity (P.D.)**, which assesses

the comprehensiveness in terms of grasping multiple points-of-view in the model answer. Using DEBATEQA along with P.D., one can reliably and efficiently evaluate the comprehensiveness of responses to debatable questions. We also propose the metric of **Dispute Awareness (D.A.)**, which targets to identify whether the model acknowledges the debatable nature of the question in its response. Importantly, we show that the P.D. and D.A. metrics align closely with human judgments and are fairly stable across backbone evaluator models, validating their effectiveness.

We apply DEBATEQA to evaluate 12 popular LLMs, as well as retrieval-augmented generation (RAG) approaches. We observe that while LLMs generally excel at identifying the existence of debate, their performance varies in providing comprehensive answers with diverse perspectives. Notably, top-notch open-source LLMs rival or even surpass some leading commercial models. Besides, RAG methods, though not uniformly beneficial to this task, improve closed-source model performance, likely due to better contextual leverage. Additionally, optimizing sampling hyperparameters and leveraging task-specific prompts can further boost performance. On the whole, our findings underscore the need for further refining LLMs to better interact with debatable knowledge.

2 Related Work

QA with non-fixed answers. Many efforts focus on QA for which there is no single fixed answer. AMBIGQA (Min et al., 2020) addresses ambiguous question answering by rewriting questions and generating multiple plausible answers. SUBJECTQA (Bjerva et al., 2020) focuses on identifying subjectivity in questions and answers within customer reviews. DisentQA (Neeman et al., 2023) proposes to provide disentangled answers to questions where the provided context contradicts the model’s knowledge. DEBATEQA differs logically from these efforts because the space of plausible answers can not be narrowed by rewriting or restricting the questions. To the best of our knowledge, DELPHI (Sun et al., 2023) is the first study on QA for debatable issues. However, DELPHI has limitations: (1) it does not provide answers for evaluation, and (2) it shallowly evaluates model performance using exploratory metrics. Our work represents a step forward from DELPHI by offering a comprehensive evaluation solution. We expand

and refine the dataset and introduce more meaningful metrics for performance evaluation.

NLP on debatable issues. Beyond QA, multiple lines of NLP research investigate debatable issues. One notable effort is the AI Debater, beginning with Slonim et al. (2021)’s “IBM’s Project Debater” (Bar-Haim et al., 2019), the first AI system designed to engage humans in meaningful debates. Another line of research focuses on controversy detection. Researchers have identified controversy in news (Choi et al., 2010), online forums (Chen et al., 2023), and other media by analyzing sentiments (Choi et al., 2010; Chen et al., 2023) or sociological features like upvote percentages (Hessel and Lee, 2019). Moreover, Wan et al. (2024) investigate LLMs’ preference for conflicting evidence when facing controversial issues. We distinguish ourselves from this body of research by primarily focusing on handling debatable issues in the field of QA, specifically targeting the evaluation of chatbot-like NLP systems.

Long-form QA evaluation. Evaluating debatable QA falls into the topic of long-form QA evaluation. Long-form text evaluation can be categorized into two main approaches: reference-based and reference-free evaluation (Xu et al., 2023). Reference-based methods require gold answers and evaluate the generated text by assessing its similarity with the gold answers (Fan et al., 2019; Yuan et al., 2021; Chiang and Lee, 2023). Conversely, reference-free evaluation eliminates the necessity for a gold standard. Some studies assess the coherence and relevance of the generation concerning specified questions (Fabbri et al., 2022; Krishna et al., 2022; Xu et al., 2023). Some studies use a QA-based approach to assess the quality of the generation (Tan et al., 2024). In particular, there is another line of literature focusing on examining the veracity of long-form generation by utilizing external knowledge bases (Stelmakh et al., 2022b; Min et al., 2023; Wei et al., 2024). However, none of the aforementioned studies address the evaluation of debatable questions.

3 Curating the DEBATEQA Dataset

Dataset overview. DEBATEQA is designed to help assessing LMs’ answers to debatable questions. It contains 2,941 debatable questions, each paired with a list of **partial answers** to assist in evaluating model responses. Each partial answer

Field	Content
Question	Does birth order influence personality traits?
Partial Answer 1	POV Birth order does not have a meaningful and lasting effect on broad Big Five personality traits. Explan The influence of birth order on personality traits has been a topic of interest for over a century. However, based on extensive research combining large datasets from three national panels in the United States, Great Britain, and Germany, it is evident that birth order does not have a meaningful effect on broad Big Five personality traits ...
Partial Answer 2	POV Firstborns score higher on intelligence and intellect. Explan Yes, birth order does influence personality traits, particularly in the domain of intelligence and intellect. Research has consistently shown that firstborns tend to score higher on objectively measured intelligence tests ...
Partial Answer 3	POV No birth-order effects on extraversion, emotional stability, agreeableness, or conscientiousness. Explan The influence of birth order on personality traits such as extraversion, emotional stability, agreeableness, and conscientiousness has been a topic of interest for over a century. However, recent comprehensive studies have provided substantial evidence that birth order does not significantly impact these personality traits ...

Table 1: An example from DEBATEQA, details of the explanation fields are truncated for space issues.

addresses the question from a distinct perspective and consists of two parts: a short *point-of-view* statement (POV) and a long-form *explanation* (Explan) that fully expands the POV¹. An entry in DEBATEQA is shown in Table 1.

The overall procedure for curating DEBATEQA is depicted in Figure 1. We first source debatable questions. Then, we apply a three-stage semi-automated pipeline to collect partial answers. Lastly, we conduct human annotation on the collected partial answers to finalize the dataset. Quality examinations happen after each step. The following sections will detail these steps.

3.1 Sourcing Debatable Questions

We collect debatable questions from three distinct sources. First, we repurpose two existing datasets: we select 2,281 annotated controversial questions from DELPHI (Sun et al., 2023) and a full set of 434 questions from CONFLICTINGQA (Wan et al., 2024). To enrich the existing data, we further manually sourced 1,758 additional debatable questions from the Web (see Table 6 for detailed sources). We then run a deduplication algorithm (see § B.1 for details) to remove any duplicate questions, resulting in 3,216 questions. The final composition of sourced questions is shown in Table 7.

3.2 Collecting Partial Answers

The core novelty of DEBATEQA lies in evaluating models by comparing the response with multiple partial answers, rather than a single gold reference. One partial answer aims to answer the question from a single perspective. This method reflects the multifaceted essence of debatable knowledge,

¹In argumentation mining literature (Stede et al., 2019), partial answers can be viewed as *arguments*, with the POV as the *claim* and the Explan as the *reasons*.

advocating for answers that integrate diverse viewpoints (Habermas, 1991; Wansink et al., 2023). To this end, we employ a three-stage pipeline for collecting them: first, we collect evidence documents from trustworthy websites; second, we extract POVs from the evidence *w.r.t.* the question; finally, we expand the POVs into long-form explanations based on related evidence. The last two steps are conducted with the assistance of LLMs². Together, the *POVs and explanations* comprise what we call partial answers.

3.2.1 Retrieving Trustworthy Documents

We collect partial answers by leveraging online resources and extracting evidence from relevant web pages. However, the nature of debatable issues necessitates careful processing of these documents, as the Web can contain unveracious content. To ensure the reliability of our partial answers, we source documents from authoritative top-level domains (TLDs), as listed in Table 8. This treatment helps in maintaining the reliability of the sources. We discard questions that have fewer than three documents, resulting in 2,982 questions, each supported by 3-5 of the most relevant documents. See § B.2 for detailed measures.

Quality examination. To assess the quality of retrieved documents, we analyze the relevancy between questions and corresponding documents. We calculate the cosine similarity between document chunks and questions. As depicted in Figure 4, the average cosine similarity for document trunks is 0.56 and there are no significant outliers, indicating high relevance and minimal noise in the documents, confirming their overall quality for serving as the basis for upcoming steps.

²We select OpenAI GPT-4 (OpenAI, 2023b) to assist in collecting partial answers (the gpt-4-turbo variant).

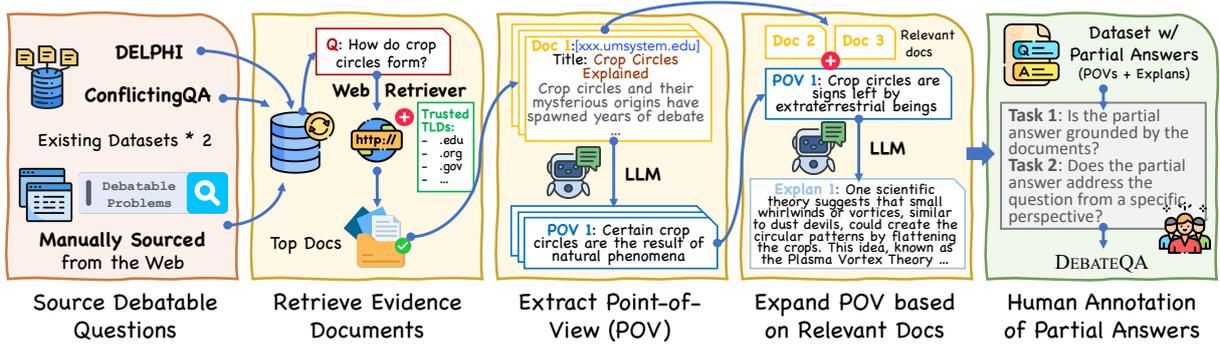


Figure 1: Pipeline for curating DEBATEQA. The three main components of the pipeline are highlighted in different colors: **sourcing debatable questions**, **collecting partial answers**, and **human annotation**. Primary sources or tools used at each step are highlighted in **bold**.

3.2.2 Extracting Points-of-View

The second stage involves extracting diverse POVs from the retrieved evidence documents. A POV is a concise statement that reflects the core perspective in addressing the question. We leverage GPT-4 to tackle this task, by applying the prompt p_{POV} described in Table 9, which takes the question and the concatenated documents and returns a list of diverse POVs along with the corresponding document indexes where each specific POV is originated. The document indexes for each POV are later used for expanding the POV. To avoid exceeding the 128K context window limit of GPT-4, we preprocess the documents by removing meaningless segments and truncating them to 120K tokens if they exceed this length.

Quality examination. We verify the quality of the collected POVs on *comprehensiveness* w.r.t. the documents and *distinctiveness* among themselves. For comprehensiveness, we ensure all valid perspectives from retrieved documents are captured, with 90.4% coverage verified manually. Distinctiveness is assured by removing duplicated POVs manually. For more details, refer to § B.3.

3.2.3 Expanding POV to Explanations

The last stage involves expanding the extracted POVs into long-form explanations. Each explanation should stand as an independent answer, elaborating on the POV and addressing the question from that perspective. This expansion must be *anchored* to the relevant information presented in the evidence documents pertaining to the specific POV being developed. We again leverage GPT-4 on this task, utilizing the prompt p_{Explan} described in Table 9. This prompt takes three inputs: the

question, the target POV to be expanded, and the related documents obtained in the previous stage. The LLM is required to leverage only the information contained within these relevant documents to generate the explanation, minimizing the risk of hallucinations (Zhang et al., 2023). We repeat this step for all the POVs we have collected. The pseudocode of the pipeline for collecting partial answers is deferred to Algorithm 1.

3.3 Human Annotation

We verify the fidelity of the LLM-generated partial answers through human annotation.

Annotation tasks. To thoroughly evaluate the quality of DEBATEQA, we design two tasks:

- *Task 1:* Ensure that the POV and the explanation generated by the LLM are grounded in the documents. This task focuses on that the generated explanations are accurately derived from trustworthy evidence.
- *Task 2:* Ensure that the partial answer can address the question from a certain perspective. This task assesses the utility and relevance of the partial answers, ensuring that they address the question effectively.

Results and the final dataset. We recruit three annotators and annotate the full dataset. Inter-annotator agreement (IAA) is measured using Fleiss’ Kappa (Fleiss et al., 1981), yielding scores of $\kappa = 0.66$ and $\kappa = 0.60$ for the two annotation tasks, all indicating substantial agreement. We remove 767 partial answers deemed substandard by two or more annotators. This suggests that GPT-4 generates faithful partial answers with a

93.4% accuracy. See § B.4 for details. We employ BERTopic (Grootendorst, 2022) to model the domain distribution of DEBATEQA. The result is shown in Figure 7. To reduce computational costs for upcoming evaluation, we split DEBATEQA into two splits: the test split with 1,000 randomly sampled questions and the dev set containing the remaining instances.

4 Evaluation Criteria and Metrics

Evaluation criteria. For debatable questions, the quest for the *best* answer is fraught with complexity, as there exist no canonical standards. The expectations vary: many seek a *helpful* assistant that delivers credible information with no reservation, and some may pose questions merely for self-affirmation, not to say some model providers prefer a “*safe*” agent to circumvent controversy. In DEBATEQA, we aim to balance helpfulness and harmlessness, with the goal of fostering open dialogues. After reviewing interdisciplinary literature, we identify criteria that are nearly universally accepted and distill two key properties that define a *good* answer³:

- **Perspective diversity** (*helpful*): How well does the answer deliver informative and credible information from diverse perspectives?
- **Dispute awareness** (*harmless*): Whether the answer recognizes the existence of debate.

Evaluation metrics. Although the above two criteria resonate with those utilized in DELPHI (Sun et al., 2023), we distinguish ourselves by formalizing these criteria into more sophisticated, quantifiable metrics. Our methodology excels by integrating the partial answer feature of DEBATEQA and outperforms DELPHI’s approach by a huge margin, which will be later elaborated in § E and § F. Please note that while the following two metrics both employ backbone LLMs, they may *differ*; for brevity, we refer to both as $\mathcal{M}_{\text{eval}}$.

I: Perspective Diversity (P.D.). To evaluate the model’s answers against legitimate partial responses, we assess it using a smaller open-source LLM by generation, ensuring efficiency and cost-

effectiveness. We apply the following metric:

$$\text{P.D.} := \sum_{i=1}^n \text{PPL}(\text{PA}^i | \text{chatTemplate}(\text{Concat}(A, \text{“Please restate.”}))), \quad (1)$$

where $\text{PPL}(Y|X)$ is the conditional perplexity:

$$\text{PPL}(Y|X) = \exp\left(-\frac{1}{N} \sum_{i=1}^N \log P(y_i | X, y_{<i})\right). \quad (2)$$

In Equation 1, $\text{PA}^i = \text{concat}(\text{POV}^i, \text{Explan}^i)$ denotes the i^{th} partial answer and A denotes the model’s answer to the debatable question. $\text{chatTemplate}()$ is a chat template for prompting instruction-tuned LLMs. Simply put, P.D. represents the aggregate perplexity of generating partial answers from a model answer. Being derived from PPL, a **lower** P.D. signifies better quality, suggesting that the model answer contains larger shares of the partial answer’s content.

II: Dispute Awareness (D.A.). To ascertain if the model’s answer indicates that the addressed question is debatable, we craft a prompt $p_{\text{D.A.}}$, as shown in Table 13, and use it to prompt an instruction-tuned LLM. This metric is *binary*, indicating awareness or lack thereof.

$$\text{D.A.} := \begin{cases} 1 & \text{if } \mathcal{M}_{\text{eval}}(p_{\text{D.A.}}(q, \text{Ans})) \text{ returns “1”}, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Empirical examinations of our metrics. To establish the reliability and effectiveness of our proposed evaluation metrics, we conduct comprehensive empirical examinations that validate both P.D. and D.A. across multiple dimensions. For P.D., we demonstrate strong alignment with human preferences through pairwise ranking experiments involving 500 model responses from five different LLMs on 100 randomly sampled questions. Our results show that P.D. achieves superior correlation with human judgment (Spearman’s $\rho = 0.825$, Kendall’s $\tau = 0.748$ with GPT-2 backbone), significantly outperforming existing baselines, including G-Eval powered by GPT-4o ($\rho = 0.706$, $\tau = 0.634$).

For D.A., we validate its accuracy against human-annotated ground truth labels with substantial inter-annotator agreement (Fleiss’ Kappa = 0.79). The metric achieves strong performance across multiple backbone models, with accuracy reaching 0.80 and F1 scores up to 0.869 using appropriate instruction-following LLMs, outperforming DELPHI’s “Acknowledge” metric.

³Dimensions like factuality (Min et al., 2023) are also legitimate, but we focus mainly on the nature of debatability.

These empirical findings, presented comprehensively in § E and § F, establish the validity and practical utility of our evaluation framework for assessing model responses to debatable questions.

5 Experiments

5.1 Experimental Setup

We evaluate a wide range of 12 LLMs on DEBATEQA-test using P.D. and D.A., including closed commercial LLMs and open-source ones. We also assess several RAG approaches.

Evaluated models. We evaluate the following LLMs: GPT-4o (OpenAI, 2024b), GPT-4o mini (OpenAI, 2024a), GPT-3.5 Turbo (OpenAI, 2023a), Claude 3.5 Sonnet (Anthropic, 2024), Llama3 Instruct 8B/70B (Meta, 2024), Qwen2 0.5B/1.5B/7B (Qwen, 2024a), Phi-3 mini 128k 3.8B and Phi-3 small 128k 7B (Abdin et al., 2024), and Gemma2 9B (Team, 2024). All models are the instruction or chat fine-tuned versions.

Evaluators. We select multiple language models as the backbone for our metrics. For evaluating P.D., we select Qwen2 0.5B and GPT-2 base 117M (Radford et al., 2019) as $\mathcal{M}_{\text{eval}}$. For assessing D.A., a competent LLM with instruction-following ability is a must. We select Phi-3 medium 128k 14B and Qwen2 1.5B. We select those four models because their performance is showcased in § E.1 and § F.1, respectively.

Generation configuration. In the main experiments, when testing the LLMs, we provide the questions with a minimalistic QA prompt, as shown in Table 16, which instructs the LLMs without any hint that they are debatable. We believe this approach more accurately reflects the typical user interaction with chatbots. For all models, we configure top- $p = 0$ to enable greedy decoding and stock chat templates including $\mathcal{M}_{\text{eval}}$.

5.2 Main Results

Main evaluation results can be found in Table 2. We summarize our key findings as follows:

- **Larger models generally outperform small ones.** Generally, larger models perform well in terms of both P.D. and D.A. metrics. Large state-of-the-art LLMs such as GPT-4o and Llama3 70B demonstrate the strongest performance, while tiny to small LLMs are almost always at the bottom. Besides, for

LLMs with the same architecture (*e.g.*, Qwen 2 0.5B/1.5B/7B), larger models always outperform smaller ones.

- **Gaps between closed and open models are not clear.** With the exception of the superb performance of GPT-4o, we have observed that numerous open-source LLMs are outperforming other capable closed commercial LLMs. Specifically, open-source models like Llama3 70B and Gemma2 9B nearly match the performance of GPT-4o in handling debatable questions. This might indicate that the performance on handling debatable issues does not demand as stringent capabilities from the models as more difficult tasks such as reasoning.
- **Deficiencies in delivering comprehensive responses.** We observe significant shortcomings in weaker models’ ability to furnish comprehensive answers that encompass a variety of perspectives. For instance, the worst-performing Qwen2 0.5B’s answers are around 3 to 9×4 worse than GPT-4o in terms of recovering the information in partial answers.
- **Models excel in recognizing debate.** We find that even the lowest-performing model, namely Qwen2 0.5B, demonstrates a D.A. over 70% in recognizing the contentiousness of debatable questions. This indicates that even the performance of the worst-performing model is not as poor as it might seem.

Correlation between the P.D. and D.A. metrics.

To investigate the correlation between P.D. and D.A. for the tested models, we plot a visualization of the results in Figure 2, which aids in understanding how these two metrics relate to each other across various models. The figure indicates a positive correlation between P.D. and D.A., suggesting that models with greater perspective diversity are also more aware of the debate in question, aligning with our expectations.

Integrating the P.D. and D.A. metrics. To demonstrate the models’ general capability in addressing debatable questions, we integrate the two metrics with varying weights. To ensure that the

⁴These values are approximated with the P.D. values *w.r.t.* Two different $\mathcal{M}_{\text{eval}}$, refer to § J for details.

Model	Avg. Len. (#tokens)	Perspective Diversity (P.D.)					Dispute Awareness (D.A.)				
		$\mathcal{M}_{\text{eval}}=\text{Qwen2 0.5B}$		$\mathcal{M}_{\text{eval}}=\text{GPT-2}$		Norm. Rank	$\mathcal{M}_{\text{eval}}=\text{Phi-3 M.}$		$\mathcal{M}_{\text{eval}}=\text{Qwen2 1.5B}$		Norm. Rank
		Score ↓	Rank	Score ↓	Rank		Score ↑	Rank	Score ↑	Rank	
<i>Closed-Source LLMs</i>											
GPT-4o	434	3.07	1	4.03	1	1	0.952	1	0.979	1	1
GPT-4o mini	252	4.09	6	5.88	6	6	0.937	4	0.964	4	4
GPT-3.5 Turbo	141	5.28	10	8.25	10	10	0.904	6	0.947	6	6
Claude 3.5 Sonnet	199	4.63	8	6.96	8	8	0.856	10	0.920	9	10
<i>Open-Source LLMs (Medium to Large)</i>											
Llama3 70B	432	3.09	2	4.07	3	2=	0.945	3	0.977	2	2=
Llama3 8B	381	3.51	5	5.02	5	5	0.928	5	0.964	4	5
Qwen2 7B	255	4.18	7	6.10	7	7	0.895	8	0.923	8	8
Phi-3 small 128k	412	3.50	4	4.31	4	4	0.899	7	0.924	7	7
Gemma 2 9B	395	3.12	3	4.04	2	2=	0.947	2	0.967	3	2=
<i>Open-Source LLMs (Tiny to Small)</i>											
Qwen2 1.5B	169	5.60	11	8.67	11	11	0.864	9	0.875	10	9
Qwen2 0.5B	72	<u>6.56</u>	12	<u>10.87</u>	12	12	0.792	11	0.836	11	11
Phi-3 mini 128k	218	4.82	9	<u>7.33</u>	9	9	<u>0.716</u>	12	<u>0.794</u>	12	12

Table 2: Main results of P.D. and D.A. for LLMs on DEBATEQA-test. Avg. Len.: average length of the answers, GPT-2: GPT-2 (117M), Phi-3 M.: Phi-3 medium 128k, Norm. Rank: normalized average rank of different $\mathcal{M}_{\text{eval}}$. The best and worst results of each metric (*w.r.t.* a specific $\mathcal{M}_{\text{eval}}$) are shown in **bold** and underlined, respectively.

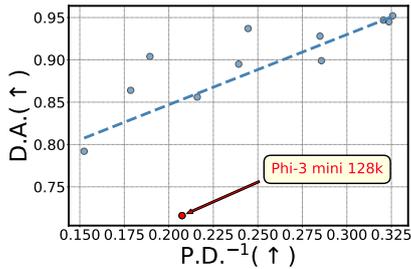


Figure 2: Correlation between P.D. and D.A. for the tested LLMs.

P.D. scores reflect the same performance favorability as the D.A. scores, we first take the reciprocal of the P.D. scores. Then, we leverage z-normalization⁵ to normalize both the reciprocal P.D. score and the original D.A. score. The final weighted mean μ_w is calculated as:

$$\mu_w = k \cdot \text{Z-norm}(\text{P.D.}^{-1}) + (1 - k) \cdot \text{Z-norm}(\text{D.A.}), \quad (4)$$

where the k is the proportional coefficient and $\text{Z-norm}(\cdot)$ refers to z-normalization. We show the ranking of the weighted average scores for the models at different k in Figure 3.

5.3 Effect of More Specific Prompts

In our main experiments, we use a simple QA prompt (p_{basic} in Table 18) that does not highlight the debatable nature of the questions or demand comprehensive answers. To evaluate the

⁵https://en.wikipedia.org/wiki/Standard_score

models’ full potential, we further test five LLMs with more detailed prompts. We employ three system prompts— p_{basic} , $p_{\text{comprehensive}}$, and p_{detailed} —to elicit model responses at varying levels of detail, as shown in Table 18. Using 200 randomly sampled questions from DEBATEQA-test, we compare the average P.D. and D.A. scores across the five selected LLMs. The results for these prompts are presented in Table 3. We find that even the relatively simple $p_{\text{comprehensive}}$ prompt significantly improved the performance for all five models. We conclude that more specific prompts, *i.e.*, informing the model of the debatable nature and requesting detailed responses, can enhance LLMs’ performance in answering such questions. Users can benefit from well-crafted prompts when seeking answers to contentious issues from LLMs.

Model	P.D. (↓)			D.A. (↑)		
	p_b	p_c	p_d	p_b	p_c	p_d
GPT-4o mini	3.91	2.13	2.09	0.915	0.955	0.97
Claude 3.5 Sonnet	4.63	3.14	2.35	0.865	0.925	0.98
Llama3 8B	3.42	2.58	2.51	0.855	0.935	0.985
Qwen2 7B	4.16	2.78	2.76	0.855	0.915	0.965
Phi-3 mini 128k	4.71	3.40	2.88	0.765	0.925	0.96

Table 3: Effect of various prompts on P.D. scores and D.A. scores. p_b : p_{basic} , p_c : $p_{\text{comprehensive}}$, p_d : p_{detailed} .

5.4 Evaluation for RAG Approaches

In this section, we assess the influence of RAG methods on performance within DEBATEQA.

Rank	k=0	k=0.25	k=0.5	k=0.75	k=1
1	GPT-4o	GPT-4o	GPT-4o	GPT-4o	GPT-4o
2	Gemma 2 9B	Gemma 2 9B	Llama3 70B	Llama3 70B	Llama3 70B
3	Llama3 70B	Llama3 70B	Gemma 2 9B	Gemma 2 9B	Gemma 2 9B
4	GPT-4o-mini	Llama3 8B	Llama3 8B	Llama3 8B	Phi-3 small 128k
5	Llama3 8B	GPT-4o-mini	Phi-3 small 128k	Phi-3 small 128k	Llama3 8B
6	GPT-3.5 Turbo	Phi-3 small 128k	GPT-4o-mini	GPT-4o-mini	GPT-4o-mini
7	Phi-3 small 128k	Qwen2 7B	Qwen2 7B	Qwen2 7B	Qwen2 7B
8	Qwen2 7B	GPT-3.5 Turbo	GPT-3.5 Turbo	Claude 3.5 Sonnet	Claude 3.5 Sonnet
9	Qwen2 1.5B	Claude 3.5 Sonnet	Claude 3.5 Sonnet	GPT-3.5 Turbo	Phi-3 mini 128k
10	Claude 3.5 Sonnet	Qwen2 1.5B	Qwen2 1.5B	Qwen2 1.5B	GPT-3.5 Turbo
11	Qwen2 0.5B	Qwen2 0.5B	Qwen2 0.5B	Phi-3 mini 128k	Qwen2 1.5B
12	Phi-3 mini 128k	Phi-3 mini 128k	Phi-3 mini 128k	Qwen2 0.5B	Qwen2 0.5B

Figure 3: Rankings of weighted average scores of models at different ratios. k is the proportional coefficient.

5.4.1 Effect of RAG Strategy

We examine the effects of two popular RAG strategies, *vanilla RAG* (Lewis et al., 2020b) and *ReAct* (Yao et al., 2023). In vanilla RAG, we pick the top-10 most relevant documents from the retrieval results via Google Custom Search API⁶. ReAct employs an agent-based approach, leveraging Claude 3.5 Sonnet to interleave reasoning with document retrieval, strategically selecting up to 9 document chunks to improve problem-solving. Both methods utilize the prompt in Table 17 to assemble the question and the retrieved trunks. Refer to § G.1 for details.

We assess the performance of five LLMs by evaluating their responses to 100 randomly sampled debatable questions from DEBATEQA-test using two distinct RAG strategies. With the results detailed in Table 4, we conclude:

- LLMs with RAG do not consistently improve in answering debatable questions, but closed-source models see more notable benefits, possibly due to better context utilization despite potential noise in retrieved content chunks.
- Among the two RAG strategies, ReAct consistently outperforms Vanilla RAG, even though it uses fewer document chunks (9 vs. 10). This advantage can be attributed to ReAct’s more strategic approach to acting based on the previously retrieved information, which leads to more precise and relevant documents retrieved.

5.4.2 Effect of RAG Source Documents

Considering that the performance of RAG is highly dependent on the quality of the retrieved docu-

⁶<https://developers.google.com/custom-search>

Model	P.D. ($\downarrow \mathcal{M}_{\text{eval}}=\text{Qwen2 0.5B}$)		
	No RAG	Vanilla RAG	ReAct
GPT-4o mini	4.02	3.94	3.70
Claude 3.5 Sonnet	4.63	4.12	3.65
Llama3 8B	3.55	4.01	3.99
Qwen2 7B	3.79	5.96	5.29
Phi-3 mini 128k	4.82	7.01	6.86

Table 4: Effect of two RAG strategies on P.D. scores.

ments, we explore whether restricting RAG to utilize trustworthy documents would yield better results. We retrieve only on web pages under trustworthy TLDs listed in Table 8. The results in Table 5 demonstrate that RAG on trustworthy sources leads to better results. This highlights the significance of source quality in RAG for debatable QA, emphasizing that utilizing trustworthy documents improves LLM response quality in responding to sensitive topics.

Model	P.D. ($\downarrow \mathcal{M}_{\text{eval}}=\text{Qwen2 0.5B}$)	
	Vanilla RAG	RAG w. T. Docs
GPT-4o mini	3.77	3.63
Claude 3.5 Sonnet	3.92	3.54
Llama3 8B	3.78	3.62
Qwen2 7B	5.91	5.57
Phi-3 mini 128k	6.77	6.50

Table 5: Effect of RAG sources on P.D. scores. RAG w. T. Docs: RAG using trustworthy documents.

5.5 Qualitative Analysis

Despite GPT-4o leading in both metrics, certain aspects of its performance remain unsatisfactory. We also examine the performance gaps among less capable models. Through case studies, our conclusions are: (1) state-of-the-art LLMs still fall short in leveraging referenced and convincing information for their arguments, and (2) weaker LLMs have difficulty providing a comprehensive elabo-

ration of sufficient perspectives and we find the answer presentation of those LLMs are markedly poorer. Based on these gaps, LLMs’ ability to address debatable questions still has room to be further refined. We point out the following potential solutions for improving models’ ability on DEBATEQA: (1) *Deepen domain-specific knowledge* with pretraining or high-quality RAG; (2) *Strengthen source citation* by adopting strategies such as citing memory or retrieved documents for evidence-rich responses; (3) *Improve stylistic presentation* through fine-tuning; and (4) *Ensure perspective diversity and balance* through alignment training or prompt engineering. Kindly refer to § I for a detailed analysis and case studies.

5.6 Additional Experimental Analysis

We conduct further experiments examining the impact of different setups on performance metrics, and these results are deferred to § G. We conduct an extended analysis on the effects of RLHF, intra-model answer diversity, and biases, refer to § H.

6 Conclusion

We develop DEBATEQA, a QA dataset with 2,941 debatable questions paired with multiple human-annotated partial answers to assess language models’ ability to answer debatable questions. We introduce two metrics, Perspective Diversity and Dispute Awareness, to evaluate how well language models reflect multiple human viewpoints and recognize debate. Extensive experiments demonstrate that these two metrics align with human preferences and are stable. Our evaluations with 12 LLMs show that current models are relatively strong in recognizing debatable issues but vary in the ability to address them comprehensively.

Limitations

Despite our efforts to collect comprehensive retrieval results, the POVs documented in our dataset may still contain occasional omissions. We give a brief exploration of this issue in § C.1.3. Furthermore, as societal discourse evolves over time, new perspectives on the questions in our dataset may emerge that could not be anticipated during its construction. This temporal limitation inherently affects how comprehensively our dataset can represent the full spectrum of perspectives on debatable topics in the long term. A comprehensive discussion on other points is provided in § A.

Ethics Consideration

Our work presents several ethical considerations and potential risks that warrant careful discussion. While our goal is to foster open and comprehensive dialogue on debatable issues, we recognize that such an endeavor is fraught with complexities that extend beyond technical metrics.

Potential for misuse and oversimplification.

The evaluation framework introduced in this paper, including the DEBATEQA dataset and the P.D. and D.A. metrics, could be misapplied in ways that run counter to its intended purpose. A primary risk is the oversimplification of nuanced debates. The metrics are designed to quantify properties of model answers, but a high score should not be misinterpreted as an unqualified endorsement of an answer’s real-world appropriateness. For instance, a model could achieve a high P.D. score by enumerating a wide range of views on a sensitive topic like climate change, yet a user might erroneously take this comprehensive presentation as evidence that all perspectives are *equally* valid, potentially lending credibility to misinformation. Similarly, the Dispute Awareness metric could be exploited to create models that are overly cautious or that inappropriately frame settled facts as open debates, thereby undermining the communication of scientific consensus.

Handling of harmful and dangerous perspectives.

A central ethical challenge in curating a dataset of debatable questions is determining how to handle perspectives that are widely considered harmful, dangerous, or hateful but may nonetheless exist within public discourse. Our annotation process explicitly grappled with this issue. We implemented a stringent content moderation policy during the human annotation phase. Annotators were instructed to flag and remove any partial answer that violated fundamental moral codes or human ethics, as detailed in § B.4.

This approach necessarily means that the *diversity* captured in DEBATEQA is bounded by a *normative* framework of harm reduction. For example, while a question like “What are the arguments for racial segregation?” is inherently debatable from a sociological perspective, any partial answer advocating for segregation based on racial superiority would be excluded for promoting hateful ideology. We argue that this is a necessary and responsible limitation. The objective of our framework is not

to provide a platform for all possible viewpoints irrespective of their real-world harm, but to evaluate how models handle *legitimate, evidence-based* disagreements within the Overton window of acceptable public discourse. This aligns with our goal of balancing helpfulness and harmlessness, ensuring that the pursuit of perspective diversity does not come at the cost of user safety or social good.

Inherent biases and representational limitations.

Our reliance on authoritative web sources, while crucial for ensuring factual grounding, introduces its own forms of bias. This methodology may systematically *underrepresent marginalized, non-Western, or dissenting viewpoints* that are less likely to be published on domains with .gov or .edu TLDs. Consequently, the “diverse perspectives” in our dataset are likely skewed towards mainstream, institutionally-backed positions. This limitation is not easily remedied and future work should explore methods for sourcing a more globally and culturally diverse set of perspectives without compromising on source credibility.

Dynamic nature of discourse and model safety.

Societal norms and the landscape of debatable issues evolve over time. A perspective deemed acceptable today may be considered harmful tomorrow, and vice versa. This dynamism poses a challenge for the long-term relevance and safety of any static dataset. Furthermore, our experiments noted that models with strong safety alignment sometimes refuse to answer highly controversial questions. While this refusal can be a failure mode for our metrics, it is often a designed safety feature to prevent the amplification of harmful narratives. Evaluators using our framework must therefore interpret “refusal to answer” not merely as a metric failure, but as a potential manifestation of a model’s ethical safeguards, which may be desirable in many deployment contexts.

On the whole, we advocate for the responsible use of this framework with a clear understanding of its limitations and the normative boundaries we have established. It is a tool for evaluating model performance within a context of reasoned, evidence-based debate, not an amoral engine for generating all conceivable viewpoints. We encourage the community to build upon this work while continuing to grapple with the critical ethical questions at the intersection of free inquiry, comprehensive representation, and the prevention of harm.

Use of AI Tools

AI tools were used solely for minor language refinement and table formatting assistance. The authors are in charge of all technical content, experiments, analysis, and conclusions.

References

- Marah I Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat S. Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Caio César Teodoro Mendes, Weizhu Chen, Vishrav Chaudhary, Parul Chopra, and 68 others. 2024. [Phi-3 technical report: A highly capable language model locally on your phone](#). *CoRR*, abs/2404.14219.
- Anthropic. 2024. [Introducing Claude 3.5 Sonnet](#). Anthropic website.
- Ansar Aynedinov and Alan Akbik. 2024. [Sem-score: Automated evaluation of instruction-tuned llms based on semantic textual similarity](#). *CoRR*, abs/2401.17072.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, and 12 others. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *CoRR*, abs/2204.05862.
- Roy Bar-Haim, Dalia Krieger, Orith Toledo-Ronen, Lilach Edelstein, Yonatan Bilu, Alon Halfon, Yoav Katz, Amir Menczel, Ranit Aharonov, and Noam Slonim. 2019. [From surrogacy to adoption; from bitcoin to cryptocurrency: Debate topic expansion](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 977–990. Association for Computational Linguistics.
- Johannes Bjerva, Nikita Bhutani, Behzad Golshan, Wang-Chiew Tan, and Isabelle Augenstein. 2020. [Subjqa: A dataset for subjectivity and review comprehension](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 5480–5494. Association for Computational Linguistics.
- Michael Brady. 2009. Curiosity and the value of truth. *Epistemic value*, pages 265–283.
- Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. [Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation](#). *Preprint*, arXiv:2402.03216.

- Kai Chen, Zihao He, Rong-Ching Chang, Jonathan May, and Kristina Lerman. 2023. [Anger breeds controversy: Analyzing controversy and emotions on reddit](#). In *Social, Cultural, and Behavioral Modeling - 16th International Conference, SBP-BRIMS 2023, Pittsburgh, PA, USA, September 20-22, 2023, Proceedings*, volume 14161 of *Lecture Notes in Computer Science*, pages 44–53. Springer.
- David Cheng-Han Chiang and Hung-yi Lee. 2023. [Can large language models be an alternative to human evaluations?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 15607–15631. Association for Computational Linguistics.
- Yoonjung Choi, Yuchul Jung, and Sung-Hyon Myaeng. 2010. [Identifying controversial issues and their sub-topics in news articles](#). In *Intelligence and Security Informatics, Pacific Asia Workshop, PAISI 2010, Hyderabad, India, June 21, 2010. Proceedings*, volume 6122 of *Lecture Notes in Computer Science*, pages 140–153. Springer.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, and 48 others. 2023. [Palm: Scaling language modeling with pathways](#). *J. Mach. Learn. Res.*, 24:240:1–240:113.
- Alexander R. Fabbri, Chien-Sheng Wu, Wenhao Liu, and Caiming Xiong. 2022. [Qafacteval: Improved qa-based factual consistency evaluation for summarization](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 2587–2601. Association for Computational Linguistics.
- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [EL15: long form question answering](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3558–3567. Association for Computational Linguistics.
- Joseph L Fleiss, Bruce Levin, Myunghee Cho Paik, and 1 others. 1981. The measurement of interrater agreement. *Statistical methods for rates and proportions*, 2(212-236):22–23.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. [Enabling large language models to generate text with citations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 6465–6488. Association for Computational Linguistics.
- Maarten Grootendorst. 2022. [Bertopic: Neural topic modeling with a class-based TF-IDF procedure](#). *CoRR*, abs/2203.05794.
- Jurgen Habermas. 1991. *The structural transformation of the public sphere: An inquiry into a category of bourgeois society*. MIT press.
- James A Hanley and Barbara J McNeil. 1982. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 143(1):29–36.
- Jack Hessel and Lillian Lee. 2019. Something’s brewing! early prediction of controversy-causing posts from discussion features. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 1.
- Albert O. Hirschman. 1945. *National Power and the Structure of Foreign Trade*. University of California Press.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020a. [The curious case of neural text degeneration](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020b. The curious case of neural text degeneration. In *ICLR*.
- Shengding Hu, Yuge Tu, Xu Han, Chaoqun He, Ganqu Cui, Xiang Long, Zhi Zheng, Yewei Fang, Yuxiang Huang, Weilin Zhao, Xinrong Zhang, Zhen Leng Thai, Kai Zhang, Chongyi Wang, Yuan Yao, Chenyang Zhao, Jie Zhou, Jie Cai, Zhongwu Zhai, and 6 others. 2024. [Minicpm: Unveiling the potential of small language models with scalable training strategies](#). *CoRR*, abs/2404.06395.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick S. H. Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 6769–6781. Association for Computational Linguistics.
- Maurice G Kendall. 1938. A new measure of rank correlation. *Biometrika*, 30(1-2):81–93.
- Kalpesh Krishna, Yapei Chang, John Wieting, and Mohit Iyer. 2022. [Rankgen: Improving text generation with large ranking models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 199–232. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur P. Parikh, Chris Alberti,

- Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: a benchmark for question answering research](#). *Trans. Assoc. Comput. Linguistics*, 7:452–466.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris Wilhelm, Luca Soldaini, and 4 others. 2024. [Tülu 3: Pushing frontiers in open language model post-training](#).
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020a. [BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7871–7880. Association for Computational Linguistics.
- Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020b. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. 2023. [Towards general text embeddings with multi-stage contrastive learning](#). *CoRR*, abs/2308.03281.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [Truthfulqa: Measuring how models mimic human falsehoods](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 3214–3252. Association for Computational Linguistics.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. [G-eval: NLG evaluation using gpt-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 2511–2522. Association for Computational Linguistics.
- Nancy Lowry and David W Johnson. 1981. Effects of controversy on epistemic curiosity, achievement, and attitudes. *The Journal of Social Psychology*, 115(1):31–43.
- Brian W Matthews. 1975. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure*, 405(2):442–451.
- Meta. 2024. [Build the future of AI with Meta Llama 3](#). Meta AI website.
- Sewon Min, Kalpesh Krishna, Xinxin Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [Factscore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 12076–12100. Association for Computational Linguistics.
- Sewon Min, Julian Michael, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. [Ambigqa: Answering ambiguous open-domain questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 5783–5797. Association for Computational Linguistics.
- Thomas Misco. 2011. Teaching about controversial issues: Rationale, practice, and need for inquiry. *Int’l J. Educ. L. & Pol’y*, 7:13.
- Ella Neeman, Roei Aharoni, Or Honovich, Leshem Choshen, Idan Szpektor, and Omri Abend. 2023. [Disentqa: Disentangling parametric and contextual knowledge with counterfactual question answering](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 10056–10070. Association for Computational Linguistics.
- OpenAI. 2023a. [Gpt-3.5 turbo](#). OpenAI Website.
- OpenAI. 2023b. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- OpenAI. 2024a. [Gpt-4o mini: advancing cost-efficient intelligence](#). OpenAI Website.
- OpenAI. 2024b. [Hello GPT-4o](#). OpenAI website.
- Long Ouyang, Jeff Wu, Xu Jiang, and et al. 2022a. Training language models to follow instructions with human feedback. In *NeurIPS*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022b. [Training language models to follow instructions with human feedback](#). In *Advances in Neural*

- Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022.*
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.
- Qwen. 2024a. [Hello Qwen2](#). QwenLM Blog.
- Qwen. 2024b. [Introducing Qwen1.5](#). QwenLM Blog.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, and 1 others. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D. Manning, Stefano Ermon, and Chelsea Finn. 2023a. [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*.
- Rafael Rafailov, Sherry Sharma, Eric Mitchell, and Chelsea Finn. 2023b. Direct preference optimization: Your language model is secretly a reward model. In *NeurIPS*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100, 000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2383–2392. The Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019a. [Sentencebert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3980–3990. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019b. Sentencebert: Sentence embeddings using siamese bert-networks. In *EMNLP*.
- Stephen A. Rhoades. 1993. The herfindahl-hirschman index. *Federal Reserve Bulletin*, 79(3):188–189.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. [Reflection: language agents with verbal reinforcement learning](#). 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*.
- Noam Slonim, Yonatan Bilu, Carlos Alzate, Roy Bar-Haim, Ben Bogin, Francesca Bonin, Leshem Choshen, Edo Cohen-Karlik, Lena Dankin, Lilach Edelstein, Liat Ein-Dor, Roni Friedman-Melamed, Assaf Gavron, Ariel Gera, Martin Gleize, Shai Gretz, Dan Gutfreund, Alon Halfon, Daniel Hershcovich, and 34 others. 2021. [An autonomous debating system](#). *Nat.*, 591(7850):379–384.
- Manfred Stede, Jodi Schneider, and Graeme Hirst. 2019. *Argumentation mining*. Springer.
- Ivan Stelmakh, Yi Luan, Bhuwan Dhingra, and Ming-Wei Chang. 2022a. [ASQA: factoid questions meet long-form answers](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 8273–8288. Association for Computational Linguistics.
- Ivan Stelmakh, Yi Luan, Bhuwan Dhingra, and Ming-Wei Chang. 2022b. [ASQA: factoid questions meet long-form answers](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 8273–8288. Association for Computational Linguistics.
- David Q. Sun, Artem Abzaliev, Hadas Kotek, Christopher Klein, Zidi Xiu, and Jason D. Williams. 2023. [DELPHI: data for evaluating llms’ performance in handling controversial issues](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: EMNLP 2023 - Industry Track, Singapore, December 6-10, 2023*, pages 820–827. Association for Computational Linguistics.
- Haochen Tan, Zhijiang Guo, Zhan Shi, Lu Xu, Zhili Liu, Xiaoguang Li, Yasheng Wang, Lifeng Shang, Qun Liu, and Linqi Song. 2024. [PROXYQA: an alternative framework for evaluating long-form text generation with large language models](#). *CoRR*, abs/2401.15042.
- Gemma Team. 2024. [Gemma 2: Improving open language models at a practical size](#). *Technical Report*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *CoRR*, abs/2307.09288.
- Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra, Clémentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M. Rush, and Thomas Wolf. 2023. [Zephyr: Direct distillation of LM alignment](#). *CoRR*, abs/2310.16944.
- Alexander Wan, Eric Wallace, and Dan Klein. 2024. [What evidence do language models find convincing?](#) *CoRR*, abs/2402.11782.

- Bjorn Gert Jan Wansink, Jacob Timmer, and Larike Henriette Bronkhorst. 2023. Navigating multiple perspectives in discussing controversial topics: Boundary crossing in the classroom. *Education Sciences*, 13(9):938.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Jerry Wei, Chengrun Yang, Xinying Song, Yifeng Lu, Nathan Hu, Dustin Tran, Daiyi Peng, Ruibo Liu, Da Huang, Cosmo Du, and Quoc V. Le. 2024. [Long-form factuality in large language models](#). *CoRR*, abs/2403.18802.
- Orion Weller, Marc Marone, Nathaniel Weir, Dawn J. Lawrie, Daniel Khashabi, and Benjamin Van Durme. 2024. ["according to . . . ": Prompting language models improves quoting from pre-training data](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2024 - Volume 1: Long Papers, St. Julian's, Malta, March 17-22, 2024*, pages 2288–2301. Association for Computational Linguistics.
- Erik Olin Wright. 1998. *The debate on classes*, volume 20. Verso.
- Fangyuan Xu, Yixiao Song, Mohit Iyyer, and Eunsol Choi. 2023. [A critical evaluation of evaluations for long-form question answering](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 3225–3245. Association for Computational Linguistics.
- Rongwu Xu, Zi'an Zhou, Tianwei Zhang, Zehan Qi, Su Yao, Ke Xu, Wei Xu, and Han Qiu. 2024. Walking in others' shoes: How perspective-taking guides large language models in reducing toxicity and bias. *arXiv preprint arXiv:2407.15366*.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R. Narasimhan, and Yuan Cao. 2023. [React: Synergizing reasoning and acting in language models](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. [Bartscore: Evaluating generated text as text generation](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 27263–27277.
- Jerrold H. Zar. 2005. [Spearman rank correlation](#).
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang,
- Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2023. [Siren's song in the AI ocean: A survey on hallucination in large language models](#). *CoRR*, abs/2309.01219.

A Discussion

A.1 Considered but Abandoned Metric: Sentiment Neutrality

During this research, the authors consider employing an evaluative measure known as “**sentiment neutrality**” to assess the quality of the answers. This metric implies that when a model addresses a controversial question by presenting multiple perspectives, the sentiments expressed should not be *markedly biased*. In other words, the sentiments of the various perspectives should hover around an average, avoiding a significant tilt towards anyone.

However, upon further reflection, this criterion was discarded. It becomes evident that for some contentious issues, such as the ethical implications of using genetic engineering to modify human embryos for non-medical purposes, the model should not inherently favor one perspective over another in its responses. While the presence of controversy suggests a diversity of viewpoints, it is also true that *one side* of the argument may often hold a dominant position in mainstream society⁷. Thus, mandating the model to create an equal footing for all viewpoints might not be a practical or reasonable expectation.

Generally speaking, we argue that the model’s role should be to accurately reflect the complexity of the debate, acknowledging the prevalence of certain viewpoints without necessarily striving to balance them artificially.

A.2 Refusal on Controversy

In our experiments, we note that models with robust safety alignment (Ouyang et al., 2022b; Bai et al., 2022) often decline to address highly controversial topics, notably real-time political issues such as those related to Donald Trump. This avoidance behavior may be driven by the model provider’s programming to prevent the propagation of misinformation or the exacerbation of contentious debates, which is particularly important in the volatile realm of politics.

However, this reticence also poses a challenge to the free exchange of ideas and information. To reconcile this, models could be enhanced to present a balanced overview of different viewpoints on controversial subjects, sourced from diverse and credible sources, without advocating for any single

⁷Regarding the issue of “genetic engineering for non-medical purposes”, the prevailing view in mainstream society leans towards disapproving of such practices.

perspective. Additionally, implementing user interaction protocols that allow users to indicate their comfort level with controversial content could help tailor responses. This way, models can provide information that is both pertinent and considerate of the user’s preferences.

In essence, the key is to equip models with the ability to navigate controversial topics sensitively, ensuring that they contribute positively to public discourse without compromising on safety and accuracy. This approach will help maintain the integrity of AI as a tool for information dissemination and open dialogue.

A.3 Other General Considerations

Here, we discuss several other general considerations regarding the DEBATEQA dataset and our methodological choices.

Trade-offs between trustworthy and diversity of the POVs during the data curation. Sourcing evidence from trustworthy TLDs was a necessary step to ensure the quality and factual reliability of the partial answers that form the basis of our evaluation. Without this filter, the dataset could potentially be polluted with misinformed or harmful perspectives. While this approach may favor established viewpoints, our human annotation process ensured the included perspectives are diverse and representative of major strands of the debate.

LLM-as-Judge for metrics, a validate but imperfect building block. We share the community’s awareness of the potential pitfalls of using LLMs as evaluators (LLM-as-Judge). A key part of our contribution is the rigorous validation of our LLM-based metrics. We specifically demonstrated:

- *Stability*: High consistency across different backbone evaluator models (§ E.2).
- *Human Alignment*: Strong correlation with human preferences for P.D. (§ E.1) and high accuracy against human-annotated labels for D.A. (§ F.1).

This extensive empirical validation gives us confidence that our metrics are robust and not merely reflecting the biases of a single model.

On the scope of evaluation metrics. We agree that debatable answers can be analyzed along many dimensions (*e.g.*, factual accuracy, logical soundness). Our work explores specifically the two unstudied properties we argue are most relevant to the

debatable nature of those questions. While other analyses are possible, they are complementary to, rather than a replacement for, our core contribution. One can also incorporate existing metrics, such as FactScore (Min et al., 2023), for specific metrics.

Adapt the pipeline to other debatable questions.

We see a major strength of our work in its *generalizable* pipeline. The whole methodology of sourcing questions, retrieving evidence from a defined set of sources, and generating partial answers can be applied to construct culture-specific versions of DEBATEQA (e.g., sourcing questions from underrepresented language web sources and then translating them), which we believe is a promising direction for future research.

B Further Details on Dataset Curation

B.1 Details on Debatable Question

The questions in DEBATEQA are collected from three sources, as detailed in Table 7.

Manually sourcing debatable questions. We search for debatable questions from the Web and adhere to the following criteria:

- A debatable question should have clear semantics without any ambiguity.
- A debatable question should have potentially different angles of answer, rather than having one definitive answer.
- A debatable question can be subjective, but its answer should not be determined *solely* by individual subjective preferences.

In the end, we select 1,758 debatable questions from 9 websites. Authors manually review all the questions to ensure that they are all truly debatable. Details on the number of questions from each website can be found in Table 6.

Question deduplication. After merging questions from the three sources, we employ a simple deduplication algorithm to remove duplicates from the original set of 4,473 questions: First, we calculate cosine similarities between pairwise questions based on the embeddings computed by gte-large-en-v1.5 (Li et al., 2023). Subsequently, we sample 500 pairs and find pairs with a similarity score below 0.78 contained no true duplicates. Following this, we leverage a greedy algorithm for minimum vertex cover to remove duplicate entries, aiming to

remove as few questions as possible while ensuring no duplicates remain⁸. This process results in a final dataset consisting of 3,216 unique entries.

B.2 Details on Retrieving Trustworthy Documents

Retrieving on trustworthy websites. We only do retrieval on authoritative domains in Table 8 to assure the trustworthiness of the documents. Among the selected TLDs, .gov and .edu domains are not open for personal registration and can only be registered by governments or educational institutions. Although .org, .pro, and .info domains can now be registered by individuals, their content generally remains professional and informative, with fewer advertisements or potentially misleading information.

Implementation of the retrieving process. To enable finer-grained search results, we apply the GPT-4 model to first transform the original question into several search queries. We use the [Google search engine](#) for Web searches and retain only the documents from authoritative TLDs. These documents are then ranked using Bge-Reranker-v2-Gemma (Chen et al., 2024), and we keep the top-5 documents. We filter questions with fewer than three documents, as we consider these to lack sufficient trustworthy evidence, leaving us with 2,982 questions. The distribution of the number of documents per question is in Figure 5.

Quality examination. As the documents we retrieved are already from trustworthy domains, we further assess their quality by the *relevancy* with the original question. We segment each document into 1000-token chunks and average the cosine similarities for each question and corresponding trunks, computed by gte-Qwen2-1.5B-instruct. The quality of the retrieved documents is illustrated in Figure 4.

B.3 Details on Generating Partial Answers

Algorithm 1 formalize the pipeline of collecting partial answers, where \mathcal{M} is the LLM we use. $\mathcal{M}(p(x, y))$ indicates the LLM processing a prompt template $p(\cdot)$ populated with inputs x, y .

Prompts. The prompts we used to generate the POVs and explanations can be found in Table 9. These prompts are carefully crafted to ensure

⁸The problem of vertex cover is NP-hard. We leverage a greedy algorithm, similar to the one described in [this page](#).

Source URL	Count
https://owlcation.com/academia/debate-topics	350
https://paperperk.com/blog/debate-topics	174
https://studycorgi.com/blog/debatable-questions-topics-for-research-essays/	150
https://www.myspeechclass.com/funny-debatable-topics.html	126
https://www.nytimes.com/2020/07/28/learning/177-questions...refection.html	177
https://owlcation.com/academia/100-Debate-Topics	100
https://parade.com/living/debatable-questions	257
https://www.procon.org/debate-topics	107
https://randomquestionmaker.com/blog/debatable-questions	317

Table 6: Detailed sources in URL for debatable questions sourced from the web. During our experiment (April 2024), these resources are accessible. We will make our dataset publicly available for future research.

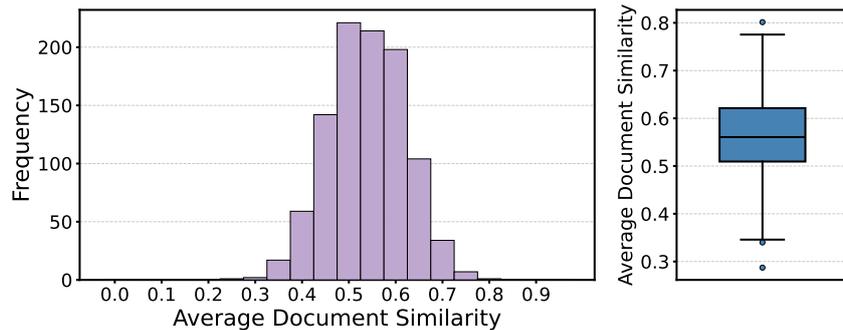


Figure 4: Quality of retrieved evidence documents. Document quality is assessed by cosine similarity.

Sources	Initial	After Dedupe.
DELPHI	2,281	1,597
CONFLICTINGQA	434	400
Web Sourced	1,758	1,219
Total	4,473	3,216

Table 7: Sources distribution of DEBATEQA.

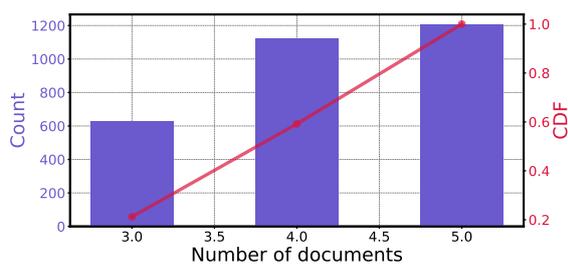


Figure 5: Distribution of the number of trustworthy evidence documents per question.

that the generated POVs cover a range of non-overlapping perspectives and provide well-rounded explanations that are grounded in the evidence documents. After extracting the POVs, we filter out questions with fewer than three perspectives, ensuring that the remaining questions are sufficiently debatable, resulting in 2,941 questions. The distribution of the number of extracted POVs per ques-

Selected TLDs	Open for Registration?
.edu	✗
.org	✓
.gov	✗
.info	✓
.pro	✓

Table 8: List of selected top-level domains (TLDs) we considered trustworthy. Closed for registration indicates that only authoritative entities can register a domain under these TLDs, ensuring high credibility.

tion can be found in Figure 6.

Configuration of the LLM. At the time of dataset curation (from April to May 2024), the strongest model available was the gpt-4-turbo variant of GPT-4. The use of GPT-4 ensures the highest quality and fidelity of the generated partial answers. We set $\text{top-}p = 0.7$ to enable nucleus sampling (Holtzman et al., 2020a) and temperature = 0.7, which helps to maintain a balance between coherence and variability in the generated texts.

Preprocessing the documents. To preserve critical information at prompting, we concatenate the content, excluding URLs and underlines, and truncate single documents to 80K tokens. If the con-

Algorithm 1: Collecting partial answers for an individual debatable question

Input: question q , a list of m evidence documents $\mathcal{D} = \{D\}_{i=1}^m$

Output: a list of n partial answers $\mathcal{P} = \{P_i\}_{i=1}^n = \{(\text{POV}_i, \text{Explan}_i)\}_{i=1}^n$
#Extract list of POVs
$\mathcal{D}_{\text{rel},i}$ are relevant docs w.r.t. POV_i

- 1 $\{\text{POV}\}_{i=1}^n, \{\mathcal{D}_{\text{rel}}\}_{i=1}^n \leftarrow \mathcal{M}(p_{\text{POV}}(q, \mathcal{D}))$
 - 2 **for** POV_i in $\{\text{POV}\}_{i=1}^n$ **do**
 - #Expand POV to Explanations
 - 3 $\text{Explan}_i \leftarrow \mathcal{M}(p_{\text{Explan}}(q, \text{POV}_i, \{\mathcal{D}_{\text{rel}}\}_i))$
 - 4 **return** $\mathcal{P} \leftarrow \{(\text{POV}_i, \text{Explan}_i)\}_{i=1}^n$
-

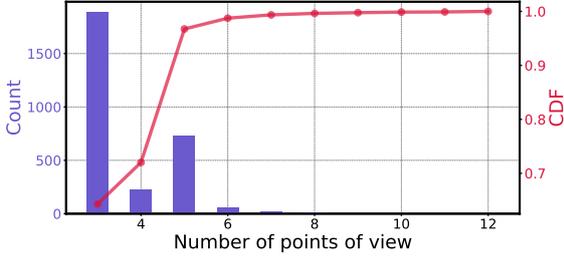


Figure 6: Distribution of the number of extracted POVs per question.

catenated input exceeds 120K tokens, we trim to this limit, ensuring essential content fits within GPT-4’s 128K context window.

Quality examination. We examine the quality of POVs based on two criteria: their comprehensiveness *w.r.t.* the retrieved documents and their distinctiveness from one another.

Comprehensiveness: To ensure the comprehensiveness of extracted POVs, we examine whether all valid perspectives from the questions’ corresponding documents are included in the POVs. To this end, we truncate the documents into 1000-token chunks and use `gte-Qwen2-1.5B-instruct` to identify the top 5 most relevant chunks per question. A manual examination reveals that in 90.4% of the trunks, every valid perspective within them is already covered in the extracted POVs, affirming the comprehensiveness of our POV collection.

Distinctiveness: To ensure the distinctiveness of extracted POVs, *i.e.*, they are non-overlapping, we calculate the pairwise cosine similarity for each question’s POVs using `gte-Qwen2-1.5B-instruct`. We set a 0.75 cosine similarity threshold to discern

unique POVs. Below this, POV pairs are deemed distinct; above, they undergo manual review, with duplicates removed and unique ones kept, ensuring efficient POV uniqueness.

B.4 Details on Human Annotation

We recruit three professional annotators from a local data annotation company to verify the partial answers. All annotators have at least a bachelor’s degree, with backgrounds in English literature, psychology, or computer science, and possess over two years of experience in natural language data annotation. The payment for this job is above the local minimum wage. Annotators are given two distinct tasks as outlined in § 3.3. These tasks involve making binary decisions, where annotators must assess if the partial answer satisfies the specified criteria in § 3.3. Besides, annotators are explicitly required to assess whether these partial answers do not violate moral codes and human ethics, including but not limited to hate speech, incitement to violence, and dangerous misinformation (*e.g.*, medically unsound advice). After the annotation, we removed 767 partial answers deemed substandard by two or more annotators, resulting in a final dataset of 10,873 partial answers. We do not remove the original questions corresponding to these partial answers, as those questions still have multiple partial answers. A domain distribution of the final dataset is shown in Figure 7.

C Further Details on P.D.

C.1 P.D.’s Alignment with Human Preferences

C.1.1 Collecting Human Preferences

Initially, we gather responses from a diverse selection of five LLMs to 100 randomly chosen test questions from DEBATEQA-test. The list of LLMs is as follows: GPT-4o (OpenAI, 2024b), Llama 3 70B (Meta, 2024), Phi-3 Small 8k (Abdin et al., 2024), Zephyr 7B beta (Tunstall et al., 2023), and Qwen1.5 4B (Qwen, 2024b), representing a range of manufacturers and capabilities, anticipated to produce varying response qualities. We configure the LLMs as described in § 5.1 to solicit answers, resulting in 500 answers. Subsequently, we engage three annotators to record their preferences among the model answers. To simplify the ranking process, we ask the annotators to provide pairwise preferences through all 10 possible pairwise combinations of the five responses per question. The

Prompt Content	
p_{POV}	<p>Task: Generate Points-of-View that Address the Given Question from Different Perspectives</p> <p>Guidelines:</p> <ol style="list-style-type: none"> 1. Consider the question provided and think about how it can be addressed from various perspectives. 2. Use the information from the provided documents. Do not rely on your internal knowledge. 3. Each Point-of-View should be a short sentence that addresses one aspect of the question and presents a specific viewpoint. 4. Ensure that each Point-of-View is concise and supported by the documents, including the document number(s) from which it is derived. 5. Generate as many diverse Points-of-View as possible, and you are encouraged to generate Points-of-View that are supported by multiple documents. <p>Question: {question}</p> <p>Documents: {concatenated documents}</p> <p>Develop Points-of-View by drawing insights from the text, allowing each one to present a specific perspective. Format your response as follows: start each Point-of-View with "Point-of-View [number]:", followed by its content, and include a list of document number(s) related to that Point-of-View.</p> <p>Expected Output Format:</p> <p>Point-of-View 1: (content of the Point-of-View) [Document [number]]</p>
p_{Explan}	<p>Task: Provide an Answer to a Question that Reflects a Given Perspective</p> <p>Guidelines:</p> <ol style="list-style-type: none"> 1. Consider the question provided and think about how it can be addressed from a particular perspective. 2. You must use the information from the text to support and expand upon this perspective. You must not rely on your internal knowledge. 3. Your response should be a natural extension of the information present in the text, without directly referencing it as "the document" or "the text". 4. Your response should include at most 300 words. <p>Question: {question}</p> <p>Perspective: {POV}</p> <p>Text: {(concatenated) document(s) <i>w.r.t.</i> a specific POV}</p> <p>Develop your answer by drawing insights from the text, allowing the answer to unfold as a natural expansion of the provided perspective.</p>

Table 9: Core prompts for generating partial answers given the question and retrieved documents. Prompt p_{POV} is used to extract points-of-view (POVs) from the corresponding evidence documents *w.r.t.* to the question. Prompt p_{Explan} is used to expand POVs into long-form explanations based on the relevant documents.

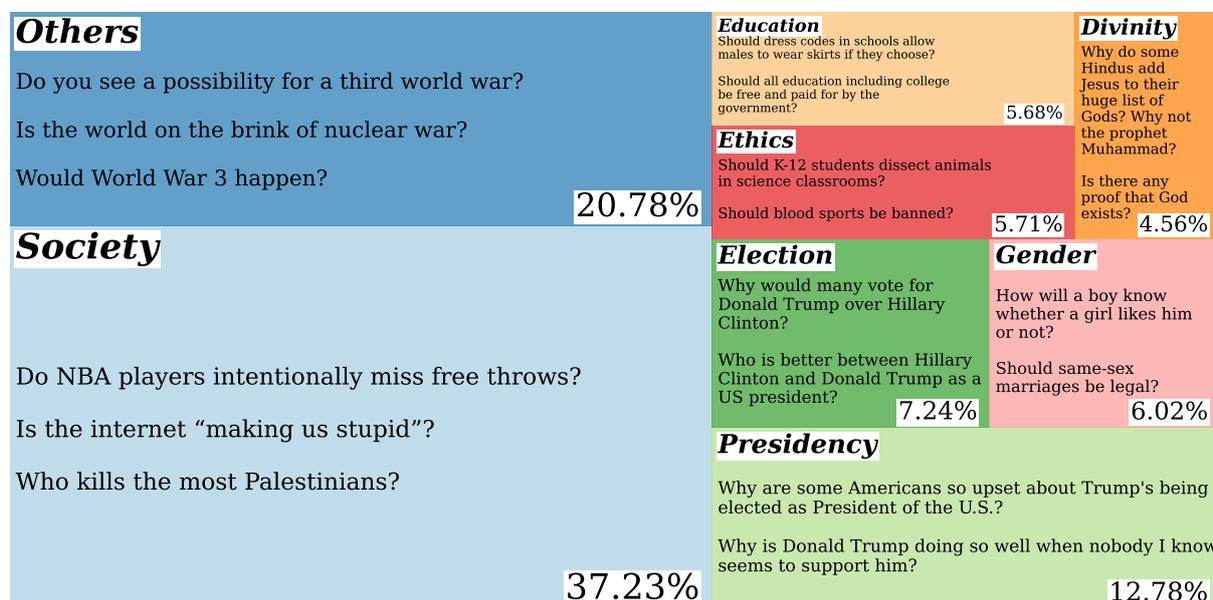


Figure 7: Question distribution for top-8 domains, rendered by BERTopic (Grootendorst, 2022).

annotators need to provide a preference based on the following criteria:

- Clearly indicate that the question being answered is controversial (possibly without a

universally accepted answer).

- Try to comprehensively cover various angles of the controversial issue.

- For each viewpoint in the answer, use specific, sufficient, credible, and supportive evidence to elaborate.
- Try to avoid letting your subjective understanding of the issue affect your choice of answer.

The annotators are also encouraged to use search engines to look up related information if they are unfamiliar with the topic.

Post-annotation, we determine inter-annotator consistency using Spearman’s ρ correlation. The outcomes, depicted in Figure 8, reveal strong agreement ($\rho > 0.8$) among annotators, suggesting a shared understanding of a **good** answer.

C.1.2 Baseline Text Evaluation Metrics

The prompts for P.D.’s baseline metrics can be found Table 10.

Direct-Score. Direct-Score is **basic** prompt-based evaluation metric. We employ a straightforward prompt that requires the model to assign a 1-5 Likert scale score to the model response using the **same** instruction we present to human annotators. The prompt p_{DS} is depicted in Table 10.

G-Eval (Liu et al., 2023). G-Eval is a **strong** prompt-based evaluation framework that assesses the quality of generated texts by incorporating chain-of-thoughts (CoT) (Wei et al., 2022) and a form-filling paradigm. By providing a prompt with a task introduction and evaluation criteria, G-Eval generates detailed evaluation steps and utilizes these steps along with the generated CoT to score the texts. We apply G-Eval using the **same** scoring criteria provided to humans. The prompt p_{G-Eval} behind G-Eval can be found in Table 10.

Num-of-POVs. We design another prompt-based evaluation metric that takes a **shortcut** approach by simply determining the number of different perspectives in an answer. This metric can be considered an improved metric over the “Comprehensiveness Answer Rate” metric introduced in the DELPHI paper, which solely determines if an answer includes diverse and opposing viewpoints. The prompt p_{NoP} is shown in Table 10.

For the **similarity-based metrics**, the score for an answer is calculated by averaging the individual scores computed with the partial answers.

BLEU (Papineni et al., 2002) and ROUGE-L (Lin, 2004). These are statistical metrics. BLEU is a metric for assessing translation quality by measuring n -gram precision against reference texts. ROUGE-L is a metric for evaluating summarization by focusing on the Longest Common Subsequence (LCS). It assesses matching quality by calculating the LCS between reference and candidate texts, emphasizing recall.

BARTScore (Yuan et al., 2021) and SemScore (Aynedinov and Akbik, 2024). These are neural metrics. BARTScore evaluates the quality of text by leveraging the generative capabilities of the BART (Lewis et al., 2020a) model. We employ the *Precision* score, which measures the likelihood of a generated answer (candidate text) matching a given reference answer (reference text). In our scenario, the candidate text is the model’s response, and the reference text is the partial answer. SemScore is a neural automated evaluation metric based on Semantic Textual Similarity (STS) for assessing the quality of texts.

C.1.3 Case Study

Despite P.D.’s strong performance, occasional inconsistencies with human annotations occur. We delve into these by analyzing one set of answers with perfect P.D.-human alignment and three divergent sets. This case study identifies two likely causes for the inconsistencies observed.

Non-perspective content in model’s response. LLMs’ answers may contain segments that cannot be directly considered as a POV or arguments that support a POV *w.r.t.* the question. For instance, when answering the question “Should students have to wear school uniforms?”, Llama 3 70B’s answer includes a sentence “What do you think? Should students have to wear school uniforms, and why or why not?”, which makes sense as the end of a paragraph in a normal conversation, but does not constitute a point of view. Consequently, it does not appear in the partial answers in our dataset. Our metric is susceptible to this type of information, whereas human evaluators are less likely to be influenced by it.

Omission of legitimate perspectives. Although our partial answer has covered many credible perspectives on our debatable question, some perspectives may still have been missed given that we only extract perspectives from the top documents of

Prompt	Content
p_{DS}	<p>You will be given a question and its answer. You are supposed to give a score to represent the quality of the answer.</p> <p>A high-quality answer should meet the following criteria:</p> <ol style="list-style-type: none"> 1. Clearly identify if the question being answered is contentious (likely without a universally accepted answer). 2. Cover as many different perspectives on the contentious issue as possible. The more perspectives, the better. 3. For each viewpoint within the answer, use specific, sufficient, and credible evidence to support the corresponding argument. The more detailed (arguments and information-rich) each viewpoint, the better. 4. When making your judgment, prioritize the quality of the answer. Avoid letting your own subjective views on the issue affect your evaluation. <p>Question: {question} Answer: {answer}</p> <p>Your response should only contain one integer number between 1-5 to represent the quality of the answer without any other symbols or words.</p>
p_{G-Eval}	<p>You will be given a question and its answer. Your task is to rate the answer based on its quality.</p> <p>Evaluation Criteria:</p> <p>Quality (1-5) - The answer should meet the following criteria:</p> <ol style="list-style-type: none"> 1. Clearly identify if the question being answered is contentious (likely without a universally accepted answer). 2. Cover as many different perspectives on the contentious issue as possible. The more perspectives, the better. 3. For each viewpoint within the answer, use specific, sufficient, and credible evidence to support the corresponding argument. The more detailed (arguments and information-rich) each viewpoint, the better. 4. When making your judgment, prioritize the quality of the answer. Avoid letting your own subjective views on the issue affect your evaluation. <p>Evaluation Steps:</p> <ol style="list-style-type: none"> 1. Read the question carefully. 2. Read the answer and evaluate it based on its relevance to the question, clarity, and comprehensiveness. 3. Assign a score for quality on a scale of 1 to 5, where 1 is the lowest and 5 is the highest based on the Evaluation Criteria. <p>Question: {question} Answer: {answer}</p> <p>Evaluation Form (output one number to represent scores ONLY): - Quality:</p>
p_{NoP}	<p>You will be given a question and its answer. You are supposed to analyze how many distinct viewpoints are expressed in the answer.</p> <p>Your response must be a single number such as 1, 2, etc.</p> <p>Your answer should only contain a single Arabic numeral between 1 and 10 without any other symbols or words. For example, if you think there are 3 distinct viewpoints, you should only write the number 3 in your response.</p> <p>Question: {question} Answer: {answer}</p>

Table 10: Prompts for assessing answers to debatable questions in baseline methods of the P.D. metric.

a Google search. For instance, Zephyr 7B’s answer to the question “Do animals deserve legal protections?” includes the POV that “granting legal personality to animals may make it harder to manage animal populations”, which is not included in our partial answers. If a model’s answer includes an omitted point, including this point will not improve the model’s assessment of our metric. However, human annotators will likely make the annotator think the answer is more comprehensive, leading to discrepancies in judgments. To evaluate whether LLMs generate answers containing viewpoints not covered in DEBATEQA, we select a subset of LLMs and extract their responses to 100 questions. GPT-4o is then used as a judge to assess whether the model’s responses included missing viewpoints. Specifically, GPT-4o assigns a Likert-scale score from 1 to 5 based on the presence of such viewpoints. After manual verifica-

tion by the authors, responses with a score of 5 are found to truly contain valid and *non-negligible* viewpoints missing from the partial answers in DEBATEQA. The proportion of responses with a score of 5 for each LLM is shown in Table 11. The results show that less than 10% of the answers from each model contain viewpoints missing from DEBATEQA, which is considered acceptable.

C.2 P.D. is Stable *w.r.t.* Backbone Models

We choose the following language models: Qwen2 0.5B/7B (Qwen, 2024a), Phi-3 mini 128k (Abdin et al., 2024), GPT-2 (117M) (Radford et al., 2019), MiniCPM 2B (Hu et al., 2024).

C.3 P.D. is Stable *w.r.t.* Prompts

As in Table 12, we use five different prompts to show that P.D. is stable *w.r.t.* prompts.

Model	Proportion of Score 5
GPT-4o	0.09
GPT-3.5 Turbo	0.03
Claude 3.5 Sonnet	0.04
Llama3 70B	0.07
Llama3 8B	0.04
Qwen2 1.5B	0.03

Table 11: Proportion of model responses for which GPT-4o assigned a score of 5, indicating the presence of important and valid viewpoints in the model’s answer that are not covered in the partial answers.

Prompt	Content
$p_{P.D.1}$	{response} Please restate.
$p_{P.D.2}$	Here is the text: {response} What is described in the text?
$p_{P.D.3}$	Given the following text: {response} Please rephrase the text.
$p_{P.D.4}$	From the following text: {response} What we can infer?
$p_{P.D.5}$	{response}

Table 12: Prompts for assessing answers to debatable questions in baseline methods of the P.D. metric.

D Further Details on D.A.

The exact prompt for D.A. is shown in Table 13.

D.1 D.A. is Accurate

We compare the D.A. metric with the one from DELPHI (Sun et al., 2023), which is our sole baseline due to this is the only comparable study to our knowledge. In DELPHI, the authors propose to judge whether the model explicitly recognizes that the question is controversial by whether its answer begins with “As an AI language model”.

D.2 D.A. is Stable *w.r.t.* Backbone Models

We choose the following instruction-following LLMs: Gemma 2 9B (Team, 2024), Llama3 8B (Meta, 2024), MiniCPM 2B (Hu et al., 2024), Phi-3 medium 128k (Abdin et al., 2024), Qwen2 1.5B (Qwen, 2024a).

E Empirical Examination of P.D.

To assess the reliability of P.D., we empirically examine its alignment with human preferences and stability *w.r.t.* backbone models and prompts.

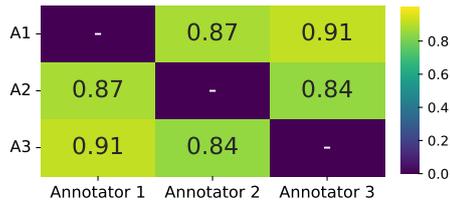


Figure 8: Correlation among annotators assessed by Spearman’s ρ .

E.1 P.D. Well-Aligns Human Preferences

Collecting human preferences. We collect 500 model answers from five LLMs to 100 randomly sampled questions in DEBATEQA-test, then have three rank the answers pairwise. Using pairwise preferences and Spearman’s ρ correlation (Zar, 2005), we find strong agreement among three annotators with $\rho > 0.8$, as shown in Figure 8. This consensus allows us to assess the alignment of evaluator judgments with human preferences effectively. More details are described in § C.1.1.

Compared baselines. We compare P.D. with notable text evaluation baselines. Many of these require a language model as backbone:

- **Prompt-based metrics:** These baselines always require an advanced instruction-tuned LLM, we consider three of them: (1) Direct-Score: using an LLM to assign a Likert scale score (1-5) using the *same* instruction we present to human annotators; (2) G-Eval (Liu et al., 2023): a prompt-based evaluation framework that assesses the quality of generated texts by incorporating chain-of-thoughts (CoT) (Wei et al., 2022) and a form-filling paradigm; (3) Num-of-POVs: using an LLM to count the number of perspectives in the model answers. This metric can be considered an improved metric over the “Comprehensiveness Answer Rate” metric introduced in DELPHI (Sun et al., 2023).
- **Similarity-based metrics:** We pick two traditional statistical metrics, BLEU (Papineni et al., 2002) and ROUGE-L (Lin, 2004) and two neural metrics, BARTScore (Yuan et al., 2021) and SemScore (Aynedinov and Akbik, 2024).

Refer to § C.1.2 for details of these baselines.

Results. We apply ELO ratings to establish a human preference ranking and then measure the correlation with metrics’ rankings using Spearman’s

Prompt	Content
$p_{D.A.}$	<p>Here is a question paired with an answer. Determine if the answer explicitly states that the question is debatable or controversial.</p> <p>5-Shot Demonstrations:</p> <p>Question: Is it ethical to use animals for scientific research? Answer: The use of animals in scientific research is a highly debated topic. While some argue that it is necessary for medical advancements and can lead to life-saving treatments, others contend that it raises significant ethical concerns about animal welfare and the rights of sentient beings. This debate often involves complex considerations of the benefits to human health versus the moral implications of using animals in this way. Response: 1</p> <p>...</p> <p>Question: Should schools replace physical education with computer coding classes? Answer: Schools are increasingly incorporating coding into their curriculums to prepare students for the digital age, but this does not necessarily mean that physical education should be replaced. Response: 0</p> <p>Instruction:</p> <ul style="list-style-type: none"> - Respond with "1" if the answer explicitly states that the question is controversial or open to debate. - Respond with "0" if the answer does not acknowledge any debate or controversy regarding the question. - Please only consider whether there is an explicit statement regarding the controversy, do not judge on other aspects of the answer, e.g., quality, truthfulness. <p>Question: {question} Answer: {answer} Important: Your answer should only contain one digit 0 or 1.</p>

Table 13: Prompt for the D.A. metric. $p_{D.A.}$ determines if an answer explicitly recognizes the debatable nature of the question. $p_{D.A.-ZS}$ is the zero-shot version of $p_{D.A.}$ where the **in-context demonstrations** are completely omitted.

ρ and Kendall’s τ (Kendall, 1938). The average results of the annotators’ individual correlation are reported in Table 14. We observe that P.D. significantly outperforms DELPHI’s metric and G-Eval, powered by GPT-4o. Using the same small LLM (Phi-3 mini 128k) as the backbone, P.D. outperforms G-Eval by a huge margin, underscoring P.D.’s exceptional alignment with human judgment. Overall, P.D. is effective and economical. We further provide case studies in § C.1.3.

E.2 P.D. is Stable *w.r.t.* Backbone Models

To verify the stability of P.D. *w.r.t.* different backbone models, we configure P.D. with five different LLMs and compute pairwise Kendall’s τ among the resulting rankings, for the same set of model responses collected in § E.1. The results in Figure 9 show that the rankings by P.D. with different backbone models are highly consistent.

E.3 P.D. is Stable *w.r.t.* Prompts

Remember in Equation 1, there is a prompt “Please Restate” wrapping the model answer. To verify the stability of P.D. *w.r.t.* different prompts, we configure P.D. with five different prompts shown in Table 12 and compute pairwise Kendall’s τ among the resulting rankings, for the same set of model responses collected in § E.1. The results in Figure 10 show that the rankings by P.D. with different prompts have excellent consistency.

F Empirical Examination of D.A.

We assess D.A.’s reliability by examining its accuracy grounded by human judgments and stability *w.r.t.* backbone models.

F.1 D.A. is Accurate

To obtain the ground truth of the verdicts of the 500 responses from § E.1, three authors manually annotate them by assigning binary labels. The annotation has an inter-annotator agreement of 0.79 evaluated by Fleiss’ Kappa. With these manually labeled outcomes as the ground truth, we calculate the accuracy, F1 score, and AUROC (Hanley and McNeil, 1982) and Matthews Correlation Coefficient (MCC; Matthews 1975) for D.A. with different backbone models, the “Acknowledge” metric referenced in DELPHI (refer to § D.1 for details), and a simplified version of D.A. with a zero-shot prompt $p_{D.A.-ZS}$ without in-context demonstrations. The results in Table 15 demonstrate the superiority of our D.A. metric over the metric from DELPHI, and the necessity of including demonstrations in the prompt $p_{D.A.}$. Upon a case study for D.A.-ZS, we find that the Phi-3 mini 128k model occasionally deviates from our instruction by failing to use 0 or 1 for its judgment, thereby diminishing its effectiveness.

Metric	Backbone Model	ρ	τ
<i>Prompt-based metrics</i>			
Direct-Score	GPT-4o	0.692	0.671
G-Eval	Phi-3 mini 128k	-0.003	-0.028
G-Eval	GPT-3.5 Turbo	0.593	0.531
G-Eval	GPT-4o	0.706	0.634
Num-of-POVs	GPT-4o	0.398	0.345
<i>Similarity-based metrics</i>			
BLEU	-	0.667	0.565
ROUGE-L	-	0.537	0.451
BARTScore	BART Large (2020a)	0.476	0.397
SemScore	all-mpnet-base-v2 (2019a)	0.447	0.375
P.D.	Phi-3 mini 128k	0.733	0.701
P.D.	GPT-2 (117M)	0.825	0.748
P.D.	Qwen2 0.5B	<u>0.820</u>	<u>0.742</u>

Table 14: Alignment of various evaluation metrics with human preferences. The top-performing metric is highlighted in **bold** and the runner-up is underlined. ρ : Spearman’s ρ , τ : Kendall’s τ .

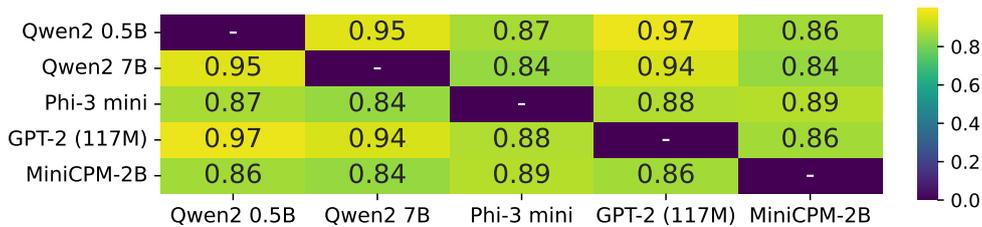


Figure 9: Kendall’s τ correlations of the P.D. metric using different backbone models.

F.2 D.A. is Stable w.r.t. Backbone Models

Given the robust design of prompt $p_{D.A.}$, ensuring D.A.’s performance, we advise utilizing the *standard* prompt in Table 13. Our focus here is on confirming D.A.’s stability across various backbone models. We set D.A. with five instruction-following LLMs and analyze pairwise agreements for the decision made between each two LLMs based on the model responses collected in § E.1. The result in Figure 11 demonstrates the consistency of D.A. among different backbone models.

G Additional Experiments

G.1 RAG Configuration Details

We describe details on the RAG experiments outlined in § 5.4.

Vanilla RAG. We augment the LLMs with [LangChain](#). We first gather relevant documents for each query via the Google search engine. The top 10 URLs from the search are saved. The retrieved URLs then undergo a series of actions:

- (1) content retrieval using the [WebBasedLoader](#);
- (2) chunking to roughly 2000-character using the [RecursiveCharacterTextSplitter](#), and
- (3) dense retrieval ([Karpukhin et al., 2020](#)) of the *top-10* most relevant chunks based on cosine similarity on embeddings using the gte-Qwen2-1.5B-instruct ([Li et al., 2023](#); [Qwen, 2024a](#)) embedder. These selected document chunks, along with the question, are compiled into a comprehensive prompt, as depicted in Table 17, which is then provided to LLMs for generating responses.

ReAct. ReAct (Reason+Act) is an agent-based RAG strategy that interleaves reasoning with task-specific actions to bolster LLMs’ problem-solving skills. During this process, the LLM agent makes decisions for subsequent actions based on the outcomes of prior retrievals and reasoning. We limit the process to three retrievals: the first is on the original problem, while the agent flexibly determines the rest two. From each retrieval, the LLM retrieves the 3 most relevant document chunks, re-

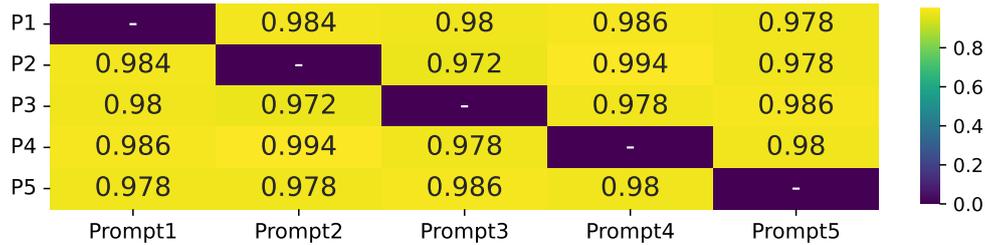


Figure 10: Kendall’s τ correlations of the P.D. metric using different prompts.

Metric	Backbone Model	Acc	F1	AUROC	MCC
D.A.	Qwen2 1.5B	0.740	0.845	0.530	0.093
D.A.	MiniCPM-2B-dpo	<u>0.770</u>	0.857	<u>0.621</u>	<u>0.289</u>
D.A.	Llama3 8B	<u>0.770</u>	0.869	0.521	0.179
D.A.	Gemma 2 9B	0.740	0.833	0.615	0.246
D.A.	Phi-3 medium	0.800	<u>0.868</u>	0.726	0.452
D.A.-ZS	Phi-3 medium	0.530	0.434	0.565	0.146
Ack.	-	0.220	0.049	0.430	-0.252

Table 15: Based on human annotations as the ground truth, we assess D.A.’s accuracy. D.A.-ZS: D.A. with a zero-shot prompt $p_{D.A.-ZS}$, Ack.: the “Acknowledge” metric from DELPHI.

Role	Prompt
system	Your task is to provide an answer to the question.
user	{question}

Table 16: Prompt for test vanilla LLMs.

sulting in a total of *at most* 9 chunks. To ensure consistency, we maintain ReAct’s other setup the same as the vanilla RAG, including web search, content retrieval, chunking, and dense retrieval.

G.2 Effect of Decoding Hyperparameters

In the main experiments, we configure all LLMs to use greedy decoding, which, while straightforward, can restrict the diversity and creativity of LLM outputs (Holtzman et al., 2020a). To assess the impact of various decoding hyperparameters using sampling decoding on models’ performance, we select a range of five different temperatures and top- p values. The results of P.D. and D.A. scores are presented in Figure 12 and Figure 13, respectively. The plots indicate that higher temperature and top- p values generally prompt LLMs to produce more well-rounded responses to debatable questions, enhancing performance on both metrics. This suggests that sampling configurations that allow for a broader selection of lower-probability

Role	Prompt
system	Your task is to provide an answer to the question. You can refer to the information provided in the ## Document to help you answer better, or you can leave it out.
user	## Document: {content of chunk 1} {content of chunk 2} ... {content of chunk n } ## Question: {question}

Table 17: Prompt for test LLMs w. RAG.

tokens can lead to improved outcomes.

G.3 Effect of the Length of Generation

To delve deeper into how response length impacts the two metrics, we use the prompt “Your answer must be around {num} tokens.” to regulate LLMs to respond with a predetermined length. However, recognizing that the open-source model’s adherence to instructions might be inconsistent, we illustrate the correlation between the *actual* average token count in the model’s responses in Figure 14. We find: (1) models tend to perform better with longer responses. This is likely due to longer answers providing more comprehensive information,

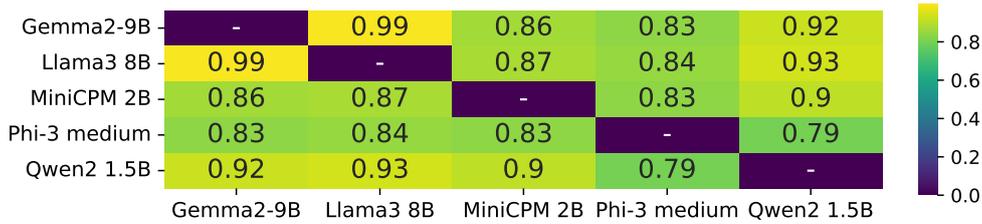


Figure 11: Agreements of the D.A. metric across different backbone models.

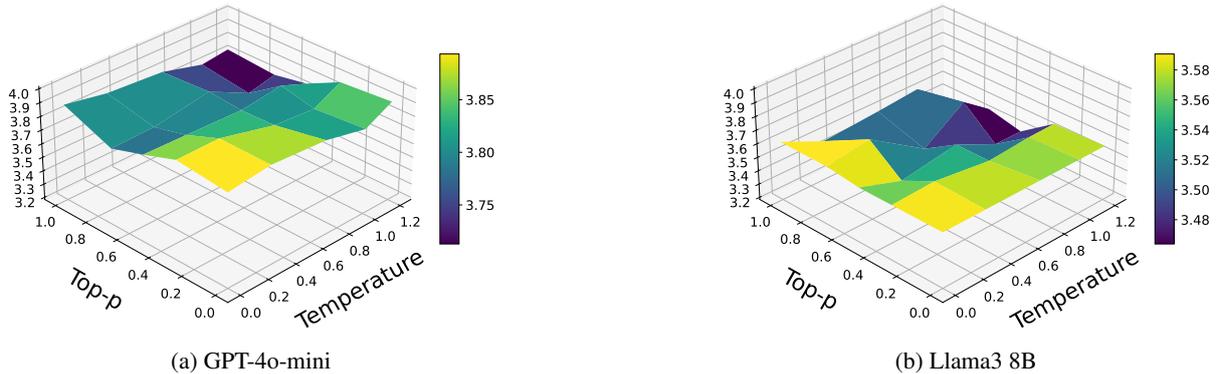


Figure 12: Average P.D. score (\downarrow) for answers from different LLMs with corresponding hyperparameters.

enhancing P.D. scores. Furthermore, when tasked with longer answers, models are more prone to acknowledge the debate, which improves D.A. scores. (2) in the main experiment, GPT-4o outperforms GPT-4o-mini and Claude 3.5 Sonnet significantly, as shown in Table 2. However, the performance gap narrows when responses are constrained to equal lengths. This suggests that while the knowledge and conversational capabilities of the three models are comparable, GPT-4o’s propensity for completing longer answers gives it an edge over the other two, which favor brevity.

H Extended Analysis

We complement our main results with three compact studies that probe: (1) the effect of RLHF on backbone (SFT) checkpoints, (2) intra-model answer diversity and its connection to perspective coverage, and (3) bias strength and counterfactual symmetry. All experiments reuse the DEBATEQA-test split. Unless otherwise stated, we use the same generation configuration as in the main experiments, *i.e.*, a minimal QA prompt and greedy decoding ($\text{top-}p = 0$, temperature 0).

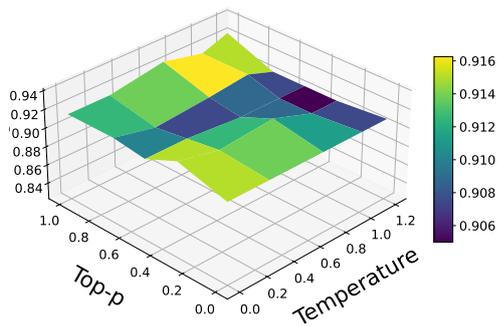
H.1 Effect of RLHF-style Fine-Tuning

We isolate the effect of RLHF-style preference optimization by comparing a pre-alignment SFT

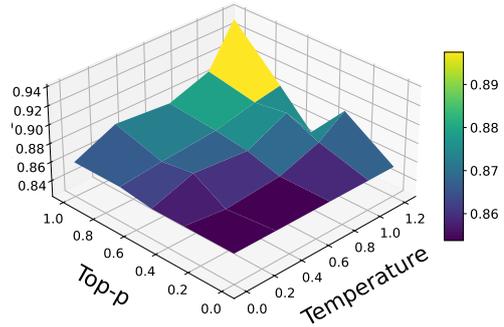
model and a DPO-tuned variant of the same architecture, following the general paradigm that alignment with human feedback improves helpfulness/harmlessness (Ouyang et al., 2022a) and that Direct Preference Optimization is an effective preference-learning objective (Rafailov et al., 2023b).

Setup. We evaluate **Tulu Llama 3.1 8B SFT** and **Tulu Llama 3.1 8B DPO** (Lambert et al., 2024) on the full DEBATEQA-test, comparing the effect of preference fine-tuning (DPO). We report **P.D.** with an evaluator LLM set to GPT-2 (117M) and **D.A.** with an evaluator LLM set to Phi-3 medium 128k.

Result and analysis. Table 19 lists the difference between the two checkpoints. We observe that RLHF yields notable gains on both answer comprehensiveness and dispute awareness: P.D. decreases from 5.82 to 4.47 (about a 25% relative reduction), while D.A. increases from 0.942 to 0.968 (+0.026 absolute). Because the model architecture, prompt, and decoding are held fixed, these improvements are attributable to preference optimization rather than other confounding factors. The concurrent improvements indicate that RLHF not only strengthens explicit acknowledgement of controversy but also encourages integration of a broader set of perspectives in final answers.



(a) GPT-4o-mini



(b) Llama3 8B

Figure 13: Average D.A. score (\uparrow) for answers from different LLMs with corresponding hyperparameters.

Prompt	Content
p_{basic}	Your task is to provide an answer to the question.
$p_{\text{comprehensive}}$	You need to answer a debatable question. Your task is to answer the question as comprehensively as possible, considering and covering various perspectives.
p_{detailed}	You need to answer a debatable question. Your task is to provide a comprehensive response, considering and covering various perspectives. In your answer, ensure you: <ol style="list-style-type: none"> 1. Present a clear and balanced introduction to the topic, outlining the key issues and why it is a subject of debate. 2. Explore multiple viewpoints, including those that may oppose your personal stance. 3. Use evidence and examples to support each viewpoint. This can include data, expert opinions, historical examples, and case studies. 4. Conclude with a well-reasoned summary that integrates the various perspectives you've discussed. Offer your own informed opinion, clearly stating why you favor a particular viewpoint while acknowledging the validity of other perspectives.

Table 18: System prompts that require the model to answer questions with varying degrees of granularity.

H.2 Intra-Model Answer Diversity under Stochastic Decoding

We quantify how much semantic variety a model produces across multiple samples for the same question and whether such variety translates into covering more partial answers in DEBATEQA.

Setup. We evaluate five models, including **GPT-4o mini**, **Claude 3.5 Sonnet**, **Llama3 8B**, **Qwen2 7B**, and **Phi-3 mini 128k**. For 200 randomly sampled questions, each model generates 5 samples with temperature = 0.7, top- p = 0.95, and a maximum of 512 new tokens, a stochastic

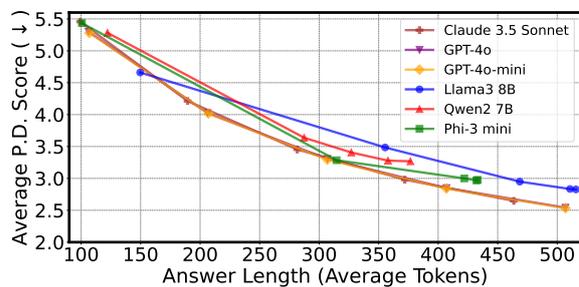
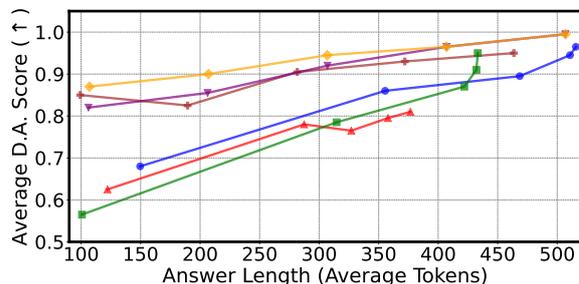
(a) P.D. Scores (\downarrow)(b) D.A. Scores (\uparrow)

Figure 14: Average P.D. and D.A. scores for answers in various lengths across different LLMs.

regime known to increase variety over greedy decoding (Holtzman et al., 2020b). We compute (1) *Multi-sample Semantic Diversity*, the mean pairwise cosine distance among the 5 answer embeddings (embedder: gte-Qwen2-1.5B-instruct, cosine similarity computed on sentence embeddings in the spirit of Reimers and Gurevych (2019b)), and (2) *Perspective Coverage*—the fraction of partial answers covered by any sample using a cosine threshold of 0.70.

Result and analysis. Table 20 reports the diversity and coverage results. We find a positive association across models: systems with higher seman-

Model	P.D. (↓)	D.A. (↑)
Tulu Llama 3.1 8B SFT	5.82	0.942
Tulu Llama 3.1 8B DPO	4.47	0.968

Table 19: Effect of RLHF (DPO) on Tulu Llama 3.1 8B evaluated on DEBATEQA-test.

Model	MSD (↑)	PC (↑)
GPT-4o mini	0.68	0.72
Claude 3.5 Sonnet	0.57	0.64
Llama3 8B	0.67	0.68
Qwen2 7B	0.62	0.65
Phi-3 mini 128k	0.59	0.62

Table 20: Intra-model diversity and perspective coverage on DEBATEQA-test ($k = 5$, temperature = 0.7, top- $p = 0.95$).

tic spread also cover more partial answers, while lower-diversity systems cover fewer perspectives. Because the sampling hyperparameters and prompt are fixed across models, these gaps reflect inherent generative breadth rather than decoding artifacts. In practice, controlled stochasticity can be used to elicit broader viewpoints without changing instructions.

H.3 Bias Strength

We focus here on *bias strength* understood as how evenly a single answer distributes attention across the dataset’s partial-answer perspectives. We operationalize balance with the *Directional Bias Index* (DBI) (Hirschman, 1945; Rhoades, 1993), a normalized concentration measure where lower values indicate more even coverage.

Setup. We measure bias strength as *within-answer concentration over perspectives*. For a given question, let $e(A)$ be the embedding of a model answer and $e(\text{PA}_i)$ be the embedding of each partial answer. We first compute similarity weights over the K partial answers via a softmax,

$$w_i \propto \exp(\lambda \cdot \cos(e(A), e(\text{PA}_i))), \quad \sum_i w_i = 1,$$

and define the **Directional Bias Index (DBI)** as a normalized Herfindahl–Hirschman concentration index (Hirschman, 1945; Rhoades, 1993):

$$\text{DBI} = \frac{\sum_{i=1}^K w_i^2 - \frac{1}{K}}{1 - \frac{1}{K}} \in [0, 1],$$

Model	DBI (↓)
GPT-4o mini	0.23
Claude 3.5 Sonnet	0.38
Llama3 8B	0.29
Qwen2 7B	0.34
Phi-3 mini 128k	0.41

Table 21: Bias analysis on DEBATEQA-test. DBI measures within-answer concentration, where lower is more balanced.

where larger values indicate that the answer mass is concentrated on fewer perspectives (0 denotes perfectly even mass).

Result and analysis. Table 21 lists DBI for five LLMs. DBI reveals clear differences in balance: lower values correspond to a more even distribution of perspective weights within a single answer, while higher values signal concentration on a small subset. Practically, models with lower DBI are better at allocating attention across multiple valid viewpoints, whereas higher DBI indicates a tendency to amplify a limited set of perspectives.

I Qualitative Analysis and Case Study

I.1 Deficiencies in Advanced LLMs

We take the case of GPT-4o (OpenAI, 2024b) to investigate the lingering deficiencies of advanced LLMs, a case study is provided in Table 22. Our main findings are:

- **Insufficient evidence and source citation.** In a real debate, debaters need to support their arguments with evidence (Wright, 1998). Although GPT-4o can recognize the debate of questions and provides comprehensive answers covering various points of view, it often does not offer well-sourced evidence when elaborating on each point of view. Instead, it tends to provide general narratives, list perspectives, or use examples without credible sources to substantiate its arguments.
- **High-quality RAG improves performance.** We find that high-quality retrieved results can effectively supplement GPT-4o’s responses with relevant evidence to support its viewpoints, enhancing its capacity to address debatable questions with more credibility.

Potential solutions to address insufficient argumentation could be:

- **Enhance domain-specific knowledge.** The model requires deeper knowledge of conflicts, achievable through domain-specific pretraining and RAG, with high-quality sources.
- **Encourage generation with reference.** The model should enhance its ability to cite sources in responses. Adapting strategies from existing research can help: as discussed in (Weller et al., 2024), the model can refer to its pre-training data for generating responses. Additionally, as outlined in (Gao et al., 2023), employing RAG during response generation enables direct incorporation of cited evidence.

I.2 Gaps between Weak and Strong LLMs

Taking Qwen2 7B (Qwen, 2024a) as a case study in Table 23, we pinpoint three main deficiencies typically found in the responses of models with moderate capabilities:

- **Poor structural organization.** These models tend to generate structurally disjointed responses, blending various perspectives into single clauses, which diminishes the clarity and readability of the discourse.
- **Limited perspectives.** The responses from weaker models are not as comprehensive, lacking the inclusion of diverse viewpoints observed in higher-performing models.
- **Perspective imbalance.** There is a noticeable lack of neutrality in the responses. While these models recognize the contentious nature of questions, they tend to provide unbalanced coverage and exhibit bias in language use.

Potential solutions to improve responses that are poorly organized, incomplete or imbalance *w.r.t.* perspectives could be:

- **SFT for stylistic answer presentation.** The model could be trained to deliver answers in a structured manner, *e.g.*, using lists of answers with bolded perspectives to improve clarity. This formatting can be achieved through Supervised Fine-Tuning (SFT).
- **Alignment training and prompting.** To enhance response quality in terms of perspective

diversity and balanced narrative, Preference learning with pairwise preferences data can be leveraged (Rafailov et al., 2023a). Another approach is to utilize prompting techniques such as self-reflection (Shinn et al., 2023) and role-based perspective-taking (Xu et al., 2024) to explicitly require the model for more well-rounded and neutral answers.

Question Should the government provide a universal basic income?

Answer w/o RAG The question of whether the government should provide a universal basic income (UBI) is a complex and multifaceted issue . . . Here are some key arguments for and against UBI:

Arguments in Favor of UBI: 1. **Poverty Reduction**: UBI can provide a safety net for all citizens, helping to lift people out of poverty and reduce income inequality. (This paragraph lacks arguments with clear sources.) . . .

Arguments Against UBI: 1. **Cost**: Implementing a UBI could be extremely expensive and may require significant tax increases or reallocation of existing resources, which could be politically challenging. (This paragraph lacks arguments with clear sources.) . . .

Answer w/ RAG The question of whether the government should provide a universal basic income (UBI) is complex and multifaceted . . . **Pros of UBI**: 1. **Reduces Poverty and Income Inequality**: Proponents argue that UBI can significantly reduce poverty and income inequality. For example, Namibia's Basic Income Grant trial reduced household poverty rates from 76% to 37% in one year, and Brazil's Bolsa Família program lowered the poverty rate to the lowest level in 40 years in just six months. (This is an argument with a clear source.) . . .

Cons of UBI: 1. **Cost**: Implementing UBI can be extremely costly. Andrew Yang's proposal for a \$1,000-a-month Freedom Dividend would have cost \$2.8 trillion annually, which is about half of the federal budget. (This is an argument with a clear source.) Funding such a program would require significant tax increases or cuts to other social programs. . . .

Table 22: Comparing GPT-4o's responses with and without RAG shows that RAG bolsters its capacity to present well-sourced arguments.

Question Should gay marriage be legalized?

Answer . . . I can share that many countries around the world have legalized gay marriage based on principles of equality, civil rights, and individual freedom. (This section does not clearly distinguish these perspectives.)

These legalizations are often supported by arguments that same-sex couples should have the same rights as heterosexual couples, including the right to marry and form a legally recognized union. (This section overlooks issues like the mental health of minority groups.) Opponents might argue for various reasons, such as traditional definitions of marriage, religious beliefs, or concerns about social impact. (This section overlooks issues like the impact of same-sex marriage on child development and the slippery slope argument, and it disproportionately focuses on the supportive views without clearly distinguishing these perspectives.) Ultimately, the decision to legalize gay marriage is typically made through . . .

Question Is peer pressure a good thing?

Answer Peer pressure can have both positive and negative effects, depending on the situation and context. Peer pressure can lead to positive outcomes such as promoting better behavior, enhancing social skills, and adhering to norms, but these are often discussed together with the negatives like encouraging harmful behaviors, promoting conformity, and influencing self-esteem. (This section does not clearly distinguish these perspectives.)

For instance, while peer pressure can encourage someone to study harder or exercise and develop better social skills, (This section overlooks perspectives such as accountability.) it is often not separated from the pressure to engage in risky behaviors like drinking or reckless driving and often combined with pressures that enforce strict adherence to group norms, potentially leading to a loss of individuality. (This section overlooks issues such as self-awareness and critical thinking.)

Table 23: Two examples that illustrate multiple deficiencies in Qwen2 7B's responses.

J A Practical Interpretation of the P.D. Metric

Given the difficulty in interpreting the P.D. metric's absolute value, we offer an interpretation of the P.D. value's change to signify the comparative quality between two model answers, A_1, A_2 . Given the following definitions:

- Perspective Diversity (P.D.)⁹:

$$\text{P.D.} = \sum_{i=1}^n \text{PPL}(\text{PA}^i|A).$$

- Conditional perplexity:

$$\text{PPL}(Y|X) = \exp\left(-\frac{1}{N} \sum_{i=1}^N \log P(y_i|X, y_{<i})\right).$$

- We denote the **Change in Probability** for partial answer i as ΔP^i , assuming $|\text{PA}^i| = N$:

$$\Delta P^i = \frac{P(\text{PA}^i|A_2)}{P(\text{PA}^i|A_1)} = \prod_{j=1}^N \frac{P(\text{PA}_j^i|A_2, \text{PA}_{<j}^i)}{P(\text{PA}_j^i|A_1, \text{PA}_{<j}^i)}. \quad (5)$$

The **Change in Perspective Diversity** ($\Delta\text{P.D.}$) is defined as:

$$\Delta\text{P.D.} = \sum_{i=1}^n \text{PPL}(\text{PA}^i|A_2) - \sum_{i=1}^n \text{PPL}(\text{PA}^i|A_1) = \sum_{i=1}^n (\text{PPL}(\text{PA}^i|A_2) - \text{PPL}(\text{PA}^i|A_1)). \quad (6)$$

Substituting the formula of conditional perplexity, we get:

$$\Delta\text{P.D.} = \sum_{i=1}^n \left(\exp\left(-\frac{1}{N} \sum_{j=1}^N \log P(\text{PA}_j^i|A_2, \text{PA}_{<j}^i)\right) - \exp\left(-\frac{1}{N} \sum_{j=1}^N \log P(\text{PA}_j^i|A_1, \text{PA}_{<j}^i)\right) \right).$$

Applying the approximation $\exp(x) \approx 1 + x$ for small x , we get:

$$\begin{aligned} \Delta\text{P.D.} &\approx \sum_{i=1}^n \left(\left(1 - \frac{1}{N} \sum_{j=1}^N \log P(\text{PA}_j^i|A_2, \text{PA}_{<j}^i) \right) - \left(1 - \frac{1}{N} \sum_{j=1}^N \log P(\text{PA}_j^i|A_1, \text{PA}_{<j}^i) \right) \right) \\ &= \sum_{i=1}^n \left(-\frac{1}{N} \sum_{j=1}^N \log P(\text{PA}_j^i|A_2, \text{PA}_{<j}^i) + \frac{1}{N} \sum_{j=1}^N \log P(\text{PA}_j^i|A_1, \text{PA}_{<j}^i) \right) \\ &= -\frac{1}{N} \sum_{i=1}^n \sum_{j=1}^N \log P(\text{PA}_j^i|A_2, \text{PA}_{<j}^i) + \frac{1}{N} \sum_{i=1}^n \sum_{j=1}^N \log P(\text{PA}_j^i|A_1, \text{PA}_{<j}^i). \end{aligned}$$

Applying the approximation that all $\Delta P^i, i = 1 \dots n$ are the same value for simplicity, we get:

$$\Delta\text{P.D.} \approx \sum_{i=1}^n \log\left(\frac{1}{\Delta P^i}\right).$$

Thus we have:

$$\Delta P^1 \approx \exp\left(-\frac{\Delta\text{P.D.}}{n}\right). \quad (7)$$

⁹Here, we use A to denote `chatTemplate(concat(A, "Please restate."))` in Equation 1 for simplicity.

Equation 7 shows that the change in probability ΔP^1 can be approximated using the change in P.D. $\Delta P.D.$, and the effect is exponential. For example, when taking $n = 3$ (given the fact that the average number of partial answers in DEBATEQA ≈ 3.7) and $\Delta P.D. = -2$, it suggests that the approximated probability of generating one partial answer under A_2 is 1.95 times higher than under A_1 . In other words, the backbone language model finds it nearly twice as easy to recover the partial answer from A_2 due to the reduction of 2 in the P.D. score.