

Coverage-based Fairness in Multi-document Summarization

Haoyuan Li¹, Yusen Zhang², Rui Zhang², Snigdha Chaturvedi¹

¹University of North Carolina at Chapel Hill, ²Pennsylvania State University
{haoyuanl, snigdha}@cs.unc.edu, {yfs5488, rmz5227}@psu.edu

Abstract

Fairness in multi-document summarization (MDS) measures whether a system can generate a summary fairly representing information from documents with different social attribute values. Fairness in MDS is crucial since a fair summary can offer readers a comprehensive view. Previous works focus on quantifying summary-level fairness using Proportional Representation, a fairness measure based on Statistical Parity. However, Proportional Representation does not consider redundancy in input documents and overlooks corpus-level unfairness. In this work, we propose a new summary-level fairness measure, **Equal Coverage**, which is based on coverage of documents with different social attribute values and considers the redundancy within documents. To detect the corpus-level unfairness, we propose a new corpus-level measure, **Coverage Parity**. Our human evaluations show that our measures align more with our definition of fairness. Using our measures, we evaluate the fairness of thirteen different LLMs. We find that Claude3-sonnet is the fairest among all evaluated LLMs. We also find that almost all LLMs overrepresent different social attribute values. The code is available at https://github.com/leehaoyuan/coverage_fairness.

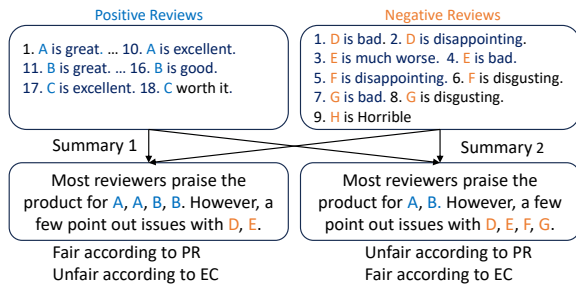
1 Introduction

Multi-document summarization (MDS) systems summarize the salient information from multiple documents about an entity, such as news articles about an event or reviews of a product. Typically, such documents are associated with *social attributes* e.g. political ideology in news and sentiments in reviews. Documents with different social attributes tend to have diverse information or conflicting opinions. A summary for them should fairly represent differing opinions across documents.


Fairness in MDS measures whether a system can generate summaries fairly representing information

from documents with different social attribute values. It can be measured at a *summary-level*—quantifying how fairly an individual summary represents documents with different social attribute values or at a *corpus-level*—quantifying how fairly a corpus of summaries as a whole represents documents with different social attribute values. Previous works in this area measured fairness in extractive or abstractive settings (Shandilya et al., 2018; Olabisi et al., 2022; Zhang et al., 2023; Huang et al., 2024). These works generally evaluate the fairness of a system as the aggregated summary-level fairness of its generated summaries. It is measured using Proportional Representation—a fairness measure based on Statistical Parity (Verma and Rubin, 2018). It requires that a document sentence being selected as a summary sentence is independent of its originating document’s social attribute value. Therefore, the distribution of social attributes among information in a fair summary should be the same as the distribution of social attributes among input documents.

The definition of Proportional Representation (PR) suffers from two key problems. The first problem is that Proportional Representation does not consider the redundancy in input documents, common in MDS. For example, in reviews of a hotel with great views, multiple reviews could comment on the view. For input documents with much redundancy, it is difficult to generate a non-redundant, fair summary according to PR. For example, suppose 67% of input reviews are positive and mostly discuss topics A and B, while 33% of reviews are negative and evenly discuss topics D, E, F, and G (Fig. 1a). According to PR, Summary 1 is fairer than Summary 2. However, Summary 2 is clearly better since, unlike Summary 1, it avoids repetition (Summary 1 repeats topic A twice) and almost equally *cover* the information from both positive and negative reviews. Hence, there is a need to redefine fairness to consider redundancy in MDS. Besides, the idea behind PR can work



(a) Existing summary-level fairness measure, Proportional Representation (PR), does not consider redundancy common in MDS. According to PR, Summary 1 is fairer than Summary 2 while Summary 2 is clearly better since it avoids repetition and cover equal proportions of information from both sentiments. Our proposed measure, Equal Coverage, correctly considers Summary 2 as fairer than Summary 1.



(b) Existing summary-level fairness measures can overlook corpus-level unfairness. Each point in this figure represents a summary sample. System 2 is fairer than System 1 since it has equal chances of overrepresenting negative (below the green line) and positive (above the green line) reviews while System 1 tends to overrepresent negative reviews. Our proposed measure, Coverage Parity, can correctly identify System 2 as fairer than System 1.

Figure 1: Issues with existing fairness measures for summary-level (a) and corpus-level (b) fairness.

for measuring fairness in other NLP tasks but not for abstractive summarization because recent LLM-based summarization methods can use quantifiers like "most" and "a few" to indicate the amount of input information *covered* by a summary sentence. Hence, there is a need to define fairness based on the proportion of input information *covered* by the summary. With this motivation, we propose a new summary-level coverage-based fairness measure, **Equal Coverage** (EC). Unlike PR, EC requires a document being *covered* by a summary sentence to be independent of its social attribute value. Since a summary sentence can cover multiple documents with similar contents, EC can address redundancy common in MDS.

The second problem is evaluating the fairness of an MDS system only using summary-level fairness measures can overlook corpus-level unfairness. Consider System 1 in Fig. 1b. Most of its summaries (blue dots) have a higher coverage probability for negative reviews than positive reviews. We observe that System 1 is unfair because its summaries *overrepresent* negative reviews. System 2 is fairer because its summaries (orange dots) have an equal chance of overrepresenting negative or positive reviews. Since individual summaries from both systems overrepresent negative or positive reviews, their summary-level fairness scores may be comparable. Hence, aggregated summary-level fairness scores cannot identify that System 2 is collectively fairer than System 1. To address this problem, we propose a new corpus-level fairness metric, **Coverage Parity**. Coverage Parity is based on the principle that documents with different social attribute

values should have equal chances of being overrepresented or underrepresented. Therefore, it can check whether the systems are equally (un)fair on different social attributes and identify which social attribute is overrepresented or underrepresented.

Our human evaluation shows that our measures align more with our definition of fairness than Proportional Representation. Using our measures, we evaluate the fairness of thirteen different Large Language Models (LLMs) (Ouyang et al., 2022; Touvron et al., 2023) in diverse domains: news, tweets, and reviews. For these domains, we consider social attributes with significant real-world impacts: ideologies and stances for news and tweets and sentiment for reviews. Our experiments find that for Llama2, Claude3, Mixtral, and Gemma2, larger models are fairer, but this trend is inconsistent for GPTs and Llama3.1. Our experiments also find that most LLMs tend to overrepresent certain social attribute values in each domain. It is an important finding that users can use to calibrate their perception before using LLM-generated summaries. It can also be used by developers to build fairer LLM-based summarization systems.

To conclude, our contributions are three-fold:

- We propose a new summary-level fairness measure, Equal Coverage, which incorporates redundancy of input information, common in MDS;
- We propose a new corpus-level fairness measure, Coverage Parity to detect corpus-level unfairness;
- We evaluate the fairness of LLMs using these two measures in various domains.

2 Related Work

Shandilya et al. (2018, 2020); Dash et al. (2019) propose to measure the summary-level fairness in MDS under the extractive setting using Proportional Representation. They propose an in-processing method to improve the fairness of extractive summaries by adding a fairness constraints to the optimization target. Similarly, Keswani and Celis (2021) uses Proportional Representation to measure the fairness of summaries under different distributions of social attributes in input documents and proposes a post-processing method to improve fairness. Olabisi et al. (2022) uses the same measure to measure fairness and proposes a clustering-based pre-processing method to improve fairness.

Recently, Zhang et al. (2023); Huang et al. (2024) extend Proportional Representation to measure fairness under the abstractive setting. To estimate the distribution of social attributes in a summary, Zhang et al. (2023) maps the summary back to the originating documents, while Huang et al. (2024) uses a finetuned model to estimate the social attribute of each summary sentence. Lei et al. (2024) measures the fairness similarly as Huang et al. (2024) and proposes to improve the fairness of abstractive summaries using reinforcement learning with a polarity distance reward. However, these works generally measure fairness using Proportional Representation, which has the limitations discussed in the introduction. In our experiments, we use Proportional Representation as a baseline.

3 Fairness Measures

In this section, we describe notation (Sec. 3.1) and our proposed measures, Equal Coverage (Sec. 3.2) and Coverage Parity (Sec. 3.3).

3.1 Notation

We use G to denote all samples for evaluating the fairness of a system on a social attribute. Each sample $(D, S) \in G$ contains a document set $D = \{d_1, \dots, d_n\}$ and a summary S generated by the system for these documents, where d_i denotes the i -th document. Each input document d_i is labeled with an social attribute value $a_i \in \{1, \dots, K\}$.

3.2 Equal Coverage

Equal Coverage is a summary-level fairness measure for measuring the fairness of a summary, S , for document sets, D . Equal Coverage is based on

the principle that the documents with different social attribute values should have equal probabilities for being covered by a summary sentence. We denote the probability that a random document $d \in D$ whose social attribute value a is k is covered by a random summary sentence $s \in S$ as $p(d, s|a = k)$. This is referred to as the coverage probability for the social attribute value k . Similarly, we denote the probability that a random document, d , is covered by a random summary sentence $s \in S$ as $p(d, s)$, which is referred to as the coverage probability for a document. For a fair summary S according to Equal Coverage, coverage of a random document d should be independent of its social attribute value:

$$p(d, s|a = k) = p(d, s), \forall k. \quad (1)$$

However, two issues arise with summaries' complex sentence structures. First, a summary sentence can combine information from several documents, making it difficult to attribute the sentence to any single document. Second, sentence lengths vary significantly based on social attribute values. The length difference can skew the coverage probability for different social attribute values $p(d, s|a = k)$. To address these issues, Equal Coverage splits the summary sentences $s \in S$ into multiple simpler sentences by prompting GPT-3.5 motivated by Bhaskar et al. (2023); Min et al. (2023). For example, compound sentences are split into simple sentences that describe a single fact. To reduce redundancy among split sentences, we filter out sentences with similar meanings based on entailment between split sentences. We denote the j -th summary sentence after split as s_j . Further details are in App. A.1.

Equal Coverage estimates the coverage probability for different social attribute values $p(d, s|a = k)$ and for a document $p(d, s)$ based on the probability $p(d_i, s_j)$ that a document d_i is covered by a summary sentence s_j . The probability $p(d_j, s_k)$ is estimated as the entailment probability that the document d_j entails the summary sentence s_k by a textual entailment model (Williams et al., 2018). Motivated by Laban et al. (2022), Equal Coverage divides documents into chunks of M words. The l -th chunk of the document d_j is denoted as $d_{j,l}$. The entailment model estimates the probability $p(d_j, s_k)$ as the maximum entailment probability $p(d_{j,l}, s_k)$ between the chunk $d_{j,l}$ and the summary sentence s_k :

$$p(d_i, s_j) = \max\{p(d_{i,l}, s_j) | d_{i,l} \in d_i\}, \quad (2)$$

where $p(d_{i,l}, s_j)$ is the probability that the document chunk $d_{i,l}$ entails the summary sentence s_j as per the entailment model. Based on the probability $p(d_i, s_j)$ that the document d_i is covered by the summary sentence s_j , the coverage probability for social attribute value i , $p(d, s | a = k)$ is empirically estimated as:

$$p(d, s | a = k) = \frac{1}{|D_k||S|} \sum_{d_i \in D_k} \sum_{s_j \in S} p(d_i, s_j), \quad (3)$$

where D_k is the set of documents d_i whose social attribute value a_i is k . Similarly, the coverage probability for a document, $p(d, s)$ is estimated as:

$$p(d, s) = \frac{1}{|D||S|} \sum_{d_i \in D} \sum_{s_j \in S} p(d_i, s_j). \quad (4)$$

Recall that for a fair summary S according to Equal Coverage, coverage probabilities for different social attribute values $p(d, s | a = k)$ should equal the coverage probability for a document $p(d, s)$. Therefore, Equal Coverage measures the summary-level fairness $EC(D, S)$ as:

$$EC(D, S) = \frac{1}{K} \sum_{k=1}^K |p(d, s) - p(d, s | a = k)|. \quad (5)$$

A high Equal Coverage value $EC(D, S)$ indicates less fairness because there are big differences between coverage probabilities for different social attribute values $p(d, s | a = k)$.

To evaluate the fairness of the system, we use the average Equal Coverage value $EC(G) = \frac{1}{|G|} \sum_{(D,S) \in G} EC(D, S)$ of all examples.

3.3 Coverage Parity

Coverage Parity is a corpus-level fairness measure designed to measure the fairness of a system of all samples G . Coverage Parity is based on the principle that documents with different social attribute values should have equal chances of being overrepresented or underrepresented by their corresponding summaries.

For a fair MDS system, Coverage Parity requires that the average difference between the coverage probability for the document d_i , $p(d_i, s)$, and the coverage probability for a document, $p(d, s)$, (Eqn. 4) is close to zero for all documents d_i whose social attribute value a_i is k . The coverage probability for the document d_i , $p(d_i, s)$ is estimated as:

$$p(d_i, s) = \frac{1}{|S|} \sum_{s_j \in S} p(d_i, s_j), \quad (6)$$

For simplicity, we denote the coverage probability difference between $p(d_i, s)$ and $p(d, s)$ as $c(d_i)$. We collect these coverage probabilities differences $c(d_i)$ from all input documents of the dataset G whose social attribute value is k into a set C_k .

Based on the average of the set of coverage probability differences C_i , Coverage Parity measures the fairness of the MDS system:

$$CP(G) = \frac{1}{K} \sum_{k=1}^K |\mathbb{E}(C_k)|, \quad (7)$$

where $\mathbb{E}(C_k)$ denotes the average coverage probability difference of C_k . A high Coverage Parity value $CP(G)$ indicates less fairness since it suggests that the chances of being overrepresented or underrepresented are very different for some social attribute values. Based on whether the average coverage probability differences, $\mathbb{E}(C_k)$, is greater or less than zero, we can also identify which social attribute value tends to be overrepresented or underrepresented.

4 Experimental Setup

We now describe experimental setup to evaluate the fairness of LLMs using our measures.

4.1 Datasets

We conduct experiments on five different datasets from the three domains: reviews, tweets, and news. These datasets are the Amazon (Ni et al., 2019), MITweet (Liu et al., 2023), SemEval (Mohammad et al., 2016), Article Bias (Baly et al., 2020), and News Stance (Ferreira and Vlachos, 2016; Pomerleau and Rao, 2017; Hanselowski et al., 2019) datasets. Tab. 1 shows the statistics of these datasets along with their social attribute values.

We observe that for some datasets, the fairness of summaries depends on the distributions of social attribute values in the input documents (Sec. 5.4). To balance social attribute values' impacts on fairness, we perform stratified sampling to collect 300 input document sets, D , for each dataset. Among these sampled sets, input document sets D dominated by different social attribute values have equal proportions. The stratified sampling does not affect the calculation of our fairness measures. More details of preprocessing are in App. A.4.

	Domain	Social Attribute	Social Attribute Values	Inp. Doc. Set Size	Avg. Doc. Len.	Avg. Doc. Len. Dif.
Amazon	Review	Sentiment	{negative, neutral, positive}	8	40	6
MITweet	Tweet	Political Ideology	{left, center, right}	20	34	3
Article Bias	News	Political Ideology	{left, center, right}	4-8	436	19
SemEval	Tweet	Stance toward Target	{support, against}	30	17	1
News Stance	News	Stance toward Target	{support, against}	4-8	240	35

Table 1: Dataset statistics for fairness evaluation in MDS. We report each dataset’s domain, social attribute, size of input document set (Inp. Doc. Set Size), average length of input documents (Avg. Doc. Len.), and average length difference between input documents with different social attribute values (Avg. Doc. Len. Dif.).

4.2 Implementation Details

We evaluate the fairness of six families of LLMs: GPTs (Ouyang et al., 2022; Achiam et al., 2023) (GPT-3.5-0124, GPT-4-turbo-2024-04-09), Llama2 (Touvron et al., 2023) (Llama2-7b, Llama2-13b, Llama2-70b), Llama3.1 (AI@Meta, 2024) (Llama3.1-8b, Llama3.1-70b), Mixtral (Jiang et al., 2023, 2024) (Mixtral-7B-Instruct-v0.1, Mixtral-8x7B-Instruct-v0.1), Gemma2 (Team et al., 2024) (gemma-2-9b-it, gemma-2-27b-it), and Claude3 (Bai et al., 2022) (claude-3-haiku-20240307, claude-3-sonnet-20240229). For comparison, we also evaluate the fairness of COOP (Iso et al., 2021) on the Amazon dataset and the fairness of PEGASUS (Zhang et al., 2020) and PRIMERA (Xiao et al., 2022) finetuned on the Multi-News (Fabbri et al., 2019) on the other datasets. The summary length is 100 words for the Article Bias and News Stance datasets and 50 for the other datasets. We describe summarization prompts in App. A.6.

To measure the difficulty of obtaining fair summaries on different datasets, we estimate lower bounds, $\text{Lower}_{\text{gre}}$, and upper bounds, $\text{Upper}_{\text{gre}}$, for both fairness measures by greedily extracting sentences that minimize or maximize the measures. We also estimate random bounds, Random , by randomly extracting sentences from input document sets as summaries. Please note that $\text{Lower}_{\text{gre}}$ and $\text{Upper}_{\text{gre}}$ are not baselines but serve as empirical bounds for our measures. The details of these bounds are in App. A.8.

To estimate the probability that a document is covered by a summary sentence (Eqn. 2), we use RoBERTa-large finetuned on the MNLI dataset (Williams et al., 2018). However, our measures are independent of this choice. We show the correlation between measures calculated using other entailment models (Schuster et al., 2021; Laurer et al., 2024) in App. A.3. To estimate the entailment probability using RoBERTa-large, we divide documents into chunks of $W = 100$ words. The

chunk size W is tuned to maximize the proportion of summary sentences whose originating documents can be identified by RoBERTa-large (App. A.2) since LLMs are less prone to factual errors (Goyal et al., 2022).

5 Experiment Results

We now describe experiment results to evaluate the fairness of LLMs using our measures.

5.1 Human Evaluation

We perform a human evaluation to determine which measure, Equal Coverage or Proportional Representation, aligns more with our definition of fairness, which is to fairly represent information from documents with different social attribute values. For Proportional Representation, we use the BARTScore (Yuan et al., 2021) implementation proposed by Zhang et al. (2023), as it shows the highest correlation with human perception.

Performing human evaluation for fairness on summarization is challenging due to the need to read entire input document sets. Therefore, we perform experiments on the Amazon dataset which only contains eight short reviews per input document set (Tab. 1). Besides, it is easier for people without training to judge whether an opinion is positive or negative than to judge if it is left-leaning or right-leaning. To simplify the evaluation, we only consider input document sets with only negative or positive reviews, displayed in two randomized columns. For each input document set, we consider the summary generated by GPT-3.5 since it shows medium-level fairness (Tab. 2, 3). To further simplify the evaluation, we focus on summaries where Equal Coverage and Proportional Representation disagree on their fairness. We randomly select 25 samples containing input document sets and corresponding summaries that meet these criteria. Each sample is annotated by three annotators recruited from Amazon Mechanical Turk. The annotators should be from English-speaking countries

	Ama.	MIT.	Art. Bia.	Sem.	New. Sta.	Ove.
Lower _{gre}	0.017	0.016	0.015	0.003	0.010	-
Upper _{gre}	0.251	0.133	0.302	0.072	0.358	-
Random	0.062	0.028	0.078	0.013	0.018	-
GPT-3.5	0.079	0.044	0.078	0.032	0.129	0.446
GPT-4	0.079	0.047	0.088	0.029	0.118	0.444
Llama2-7b	0.088	0.046	0.098	0.032	0.132	0.712
Llama2-13b	0.074	0.047	0.094	0.033	0.116	0.501
Llama2-70b	0.074	0.048	0.108	0.030	0.118	0.520
Llama3.1-8b	0.081	0.043	0.094	0.026	0.123	0.304
Llama3.1-70b	0.077	0.044	0.095	0.030	0.115	0.341
Mistral-7b	0.083	0.048	0.111	0.032	0.130	0.773
Mixtral-8x7b	0.068	0.043	0.127	0.029	0.119	0.356
Gemma2-9b	0.084	0.043	0.091	0.026	0.112	0.212
Gemma2-27b	0.075	0.043	0.091	0.026	0.116	0.182
Claude3-haiku	0.079	0.047	0.085	0.034	0.142	0.709
Claude3-sonnet	0.077	0.043	0.087	0.025	0.128	0.251
COOP	0.146	-	-	-	-	-
PEGASUS	-	0.052	0.126	0.027	0.134	-
PRIMERA	-	0.056	0.131	0.025	0.137	-

Table 2: Average summary-level fairness and overall scores on different datasets according to Equal Coverage. A lower value indicates a fairer system. **Bold** indicates the fairest system.

and have HIT Approval Rates greater than 98%. More details are in App. A.5.

Following previous practice of performing human evaluation on fairness in summarization (Shandilya et al., 2020), annotators are asked to read all reviews and identify all unique opinions. They then are asked to read the summary and rate its fairness as leaning negative, fair, or leaning positive. The Randolph’s Kappa (Randolph, 2005) between annotations of three annotators is 0.42, which shows moderate correlation. Out of these 25 samples, annotations aligns more with Equal Coverage in 17 samples, while aligns more with Proportional Representation in 8 samples. The difference is statistically significant ($p < 0.05$) using paired bootstrap resampling (Koehn, 2004). It shows that our implementation aligns more with our definition of fairness than Proportional Representation.

We compare Coverage Parity with second-order fairness (Zhang et al., 2023) based on Proportional Representation to evaluate their effectiveness in identifying overrepresented sentiments on the group level. For this, we perform bootstrapping to generate 5000 groups of bootstrap samples. We find that Coverage Parity aligns with human annotation in 95% of the groups, while second-order fairness aligns in 5% of the groups. It shows that Coverage Parity aligns more with our definition of fairness than the second-order fairness.

5.2 Summary-level Fairness Evaluation

To evaluate the summary-level fairness of different LLMs, we report the average Equal Coverage values for all samples in each dataset, $EC(G)$. We also report an *Overall* score which is the average of normalized $EC(G)$ ($[0, 1]$) using min-max normalization on all datasets. Results are in Tab. 2.

From the table, we observe that Gemma2-27b is the fairest based on the Overall score. Among smaller LLMs (with 7 to 9 billions parameters), Gemma2-9b is the fairest. We also observe that almost all evaluated LLMs are fairer than COOP on the Amazon dataset, and PEGASUS and PRIMERA on the MITweet and Article Bias datasets. For the comparison within families of GPTs, Llama2, Claude3, Mixtral, and Gemma2, we observe that larger models are generally fairer.

Previous works by Zhang et al. (2023) and Lei et al. (2024) evaluate the fairness of LLMs using Proportional Representation. When evaluating the fairness on sentiments in the review domain, our results using Equal Coverage are consistent with these works in the finding that GPT-4 is fairer than GPT-3.5, and both are fairer than COOP. However, they find that Llama2-13b is the fairest, while our results show that Llama2-13b and Llama2-70b are comparably fair. When evaluating the fairness on political ideologies in the tweet domain, our results are consistent with these works on that smaller models are fairer for Llama2. However, they find that GPT-4 is fairer than GPT-3.5 while we find that GPT-3.5 is fairer than GPT-4. We also show that GPT-4 is fairer than Llama2-7b, which contrasts the previous works. We additionally report the results for Proportional Representation on our datasets in App. A.9 and also find much inconsistency between Equal Coverage and Proportional Representation. As shown in Sec. 1, Equal Coverage considers redundancy common in MDS, suggesting better reflects fairness in MDS.

5.3 Corpus-level Fairness Evaluation

To evaluate the corpus-level fairness of different LLMs, we report the Coverage Parity, $CP(G)$, on different datasets. For each dataset, we report the most overrepresented and underrepresented social attribute value k whose average coverage probability difference, $\mathbb{E}(C_k)$, is the maximum and minimum respectively. We also report an *Overall* score which is the average of normalized $CP(G)$ ($[0, 1]$) using min-max normalization on all datasets. The

	Ama.			MIT.			Art. Bia.			Sem.			New. Sta.			Ove.
	$CP(G)$	over	under	$CP(G)$	over	under	$CP(G)$	over	under	$CP(G)$	over	under	$CP(G)$	over	under	
Lower _{gre}	0.000	neu.	pos.	0.000	lef	cen.	0.000	rig.	lef.	0.000	aga.	sup.	0.000	sup.	aga.	-
Upper _{gre}	0.141	<u>neg.</u>	<u>pos.</u>	0.048	<u>rig.</u>	<u>lef.</u>	0.195	<u>rig.</u>	<u>lef.</u>	0.040	<u>sup.</u>	<u>aga.</u>	0.256	<u>aga.</u>	<u>sup.</u>	-
Random	0.008	neg.	neu.	0.003	rig.	cen.	0.005	rig.	cen.	0.001	sup.	aga.	0.001	sup.	aga.	-
GPT-3.5	0.019	<u>neg.</u>	<u>pos.</u>	0.006	<u>lef.</u>	<u>cen.</u>	0.011	cen.	<u>lef.</u>	0.009	<u>sup.</u>	<u>aga.</u>	0.016	<u>aga.</u>	<u>sup.</u>	0.598
GPT-4	0.020	<u>neg.</u>	<u>pos.</u>	0.004	<u>lef.</u>	<u>cen.</u>	0.015	rig.	<u>lef.</u>	0.008	<u>sup.</u>	<u>aga.</u>	0.033	<u>aga.</u>	<u>sup.</u>	0.622
Llama2-7b	0.022	<u>neg.</u>	<u>pos.</u>	0.006	<u>lef.</u>	<u>cen.</u>	0.017	<u>rig.</u>	<u>lef.</u>	0.008	<u>sup.</u>	<u>aga.</u>	0.011	aga.	sup.	0.634
Llama2-13b	0.010	<u>neg.</u>	<u>pos.</u>	0.004	<u>lef.</u>	cen.	0.015	<u>rig.</u>	<u>lef.</u>	0.011	<u>sup.</u>	<u>aga.</u>	0.005	aga.	sup.	0.399
Llama2-70b	0.008	<u>neg.</u>	<u>neu.</u>	0.003	<u>lef.</u>	cen.	0.015	<u>rig.</u>	<u>lef.</u>	0.007	<u>sup.</u>	<u>aga.</u>	0.002	aga.	sup.	0.182
Llama3.1-8b	0.021	<u>neg.</u>	<u>pos.</u>	0.004	<u>lef.</u>	<u>cen.</u>	0.013	<u>rig.</u>	<u>lef.</u>	0.008	<u>sup.</u>	<u>aga.</u>	0.006	aga.	sup.	0.442
Llama3.1-70b	0.016	<u>neg.</u>	<u>pos.</u>	0.006	<u>lef.</u>	<u>cen.</u>	0.016	<u>rig.</u>	<u>lef.</u>	0.004	<u>sup.</u>	<u>aga.</u>	0.004	aga.	sup.	0.436
Mistral-7b	0.016	<u>neg.</u>	<u>pos.</u>	0.005	<u>lef.</u>	<u>cen.</u>	0.017	<u>rig.</u>	<u>lef.</u>	0.010	<u>sup.</u>	<u>aga.</u>	0.015	<u>aga.</u>	<u>sup.</u>	0.574
Mistral-8x7b	0.011	<u>neg.</u>	neu.	0.006	<u>lef.</u>	<u>cen.</u>	0.023	<u>rig.</u>	<u>lef.</u>	0.002	<u>sup.</u>	<u>aga.</u>	0.020	<u>aga.</u>	<u>sup.</u>	0.494
Gemma2-9b	0.022	<u>neg.</u>	<u>pos.</u>	0.006	<u>lef.</u>	<u>cen.</u>	0.014	cen.	<u>lef.</u>	0.008	<u>sup.</u>	<u>aga.</u>	0.012	<u>aga.</u>	<u>sup.</u>	0.635
Gemma2-27b	0.021	<u>neg.</u>	<u>pos.</u>	0.007	<u>lef.</u>	<u>cen.</u>	0.010	<u>rig.</u>	<u>lef.</u>	0.006	<u>sup.</u>	<u>aga.</u>	0.010	aga.	sup.	0.517
Claude3-hai.	0.016	<u>neg.</u>	<u>pos.</u>	0.004	<u>lef.</u>	<u>cen.</u>	0.010	cen.	<u>lef.</u>	0.010	<u>sup.</u>	<u>aga.</u>	0.037	<u>aga.</u>	<u>sup.</u>	0.531
Claude3-son.	0.017	<u>neg.</u>	<u>pos.</u>	0.004	<u>lef.</u>	<u>cen.</u>	0.012	<u>rig.</u>	<u>lef.</u>	0.002	<u>sup.</u>	<u>aga.</u>	0.025	<u>aga.</u>	<u>sup.</u>	0.313
COOP	0.048	<u>pos.</u>	<u>neg.</u>	-	-	-	-	-	-	-	-	-	-	-	-	-
PEGASUS	-	-	-	0.001	lef	cen.	0.009	rig.	lef.	0.002	sup.	aga.	0.012	sup.	aga.	-
PRIMERA	-	-	-	0.002	rig.	cen.	0.020	<u>rig.</u>	<u>lef.</u>	0.000	sup.	aga.	0.002	sup.	aga.	-

Table 3: Coverage Parity, $CP(G)$, and the most overrepresented (over) and underrepresented (under) social attribute values, and overall scores on different datasets. A lower value of $CP(G)$ indicates a fairer system. **Bold** indicates the fairest system. The social attribute values whose average coverage probability differences are statistically significantly ($p < 0.05$) different from zero based on Bootstrapping are underlined. Most LLMs overrepresent and underrepresent certain social attribute values on different datasets

results are in Tab. 3.

From the table, we observe that Llama2-70b is the fairest based on the Overall score. Among smaller LLMs (with 7 to 9 billions parameters), Llama3.1-8b is the fairest. We also observe that evaluated LLMs are less fair than PEGASUS and PRIMERA on the SemEval dataset. While comparing within each family of LLMs, we observe that larger models are fairer for the families of Llama2, Llama3.1, Mistral, Gemma2, and Claude3. However, for the family of GPTs, GPT-4 is less fair than GPT-3.5, suggesting that larger models are not necessarily more fair on the corpus level. Besides, we observe that the fairness measured by Coverage Parity and Equal Coverage are different on some datasets. The difference indicates that we should consider both summary-level fairness and corpus-level fairness for comprehensively measuring fairness in MDS.

We can also observe that most LLMs overrepresent and underrepresent certain social attribute values on different datasets. For the Amazon dataset, most LLMs overrepresent negative reviews and underrepresent positive reviews. Contrarily, COOP overrepresents positive reviews and underrepresents negative reviews. For the MITtweet and Article Bias datasets, all LLMs overrepresent left-leaning documents in the tweet domain, while most LLMs overrepresent right-leaning documents in the news domain. Contrarily, PEGASUS and

PRIMERA overrepresent right-leaning documents for all domains. We can observe the same pattern for the SemEval and News Stance datasets. All LLMs overrepresent documents supporting the target in the tweet domain but overrepresent documents against the target in the news domain. These results indicate that summaries generated by LLMs might overrepresent documents with certain social attribute values. Users should know this before they make judgments based on these summaries. For example, users should know that a review summary generated by LLMs for a product might unfairly overrepresent negative reviews so they can calibrate their perception of the product accordingly. Developers can also build fairer LLMs for MDS based on these results.

5.4 Fairness under Different Distributions of Social Attributes

We perform experiments to evaluate whether the fairness of LLMs changes under different distributions of social attribute values in input document sets. For this, we divide all samples G into K non-overlapping sets: $\{G_1, \dots, G_K\}$ based on distributions of social attribute values. Each set G_k includes samples where most documents $d \in D$ have a social attribute value of k . We denote the set G_k as the set dominated by social attribute value k . To measure differences of the summary-level fairness measured by Equal Coverage under differ-

	Ama.	MIT.	Art. Bia.	Sem.	New. Sta.
GPT-3.5	<u>0.023</u>	0.004	0.005	0.004	<u>0.022</u>
GPT-4	<u>0.023</u>	0.009	<u>0.015</u>	<u>0.007</u>	<u>0.018</u>
Llama2-7b	<u>0.023</u>	<u>0.008</u>	<u>0.018</u>	0.004	<u>0.014</u>
Llama2-13b	<u>0.012</u>	0.006	0.010	0.004	0.004
Llama2-70b	<u>0.020</u>	0.001	0.009	0.003	0.010
Llama3.1-8b	<u>0.021</u>	<u>0.005</u>	<u>0.008</u>	0.002	<u>0.031</u>
Llama3.1-70b	<u>0.019</u>	<u>0.007</u>	0.010	0.003	0.020
Mistral-7b	<u>0.020</u>	<u>0.005</u>	<u>0.018</u>	<u>0.001</u>	<u>0.012</u>
Mixtral-8x7b	<u>0.010</u>	0.003	0.003	0.004	0.012
Gemma2-9b	<u>0.017</u>	<u>0.007</u>	<u>0.010</u>	0.003	<u>0.020</u>
Gemma2-27b	<u>0.019</u>	0.003	0.003	0.001	0.004
Claude3-hai.	<u>0.015</u>	0.006	0.004	0.001	0.008
Claude3-son.	<u>0.016</u>	0.003	0.060	0.004	<u>0.028</u>

Table 4: Differences of Equal Coverage under different distributions of social attribute values in input document sets. Statistically significant ($p < 0.05$) differences based on Bootstrapping are underlined. We observe significant differences in summary-level fairness for most LLMs on the Amazon dataset.

ent distributions, we use maximum differences of equal coverage values $EC(G_k)$ (Sec. 3.2) on sets dominated by different social attribute values G_k . The results are shown in Tab. 4. From the table, we can observe that summary-level fairness changes much on the Amazon dataset but is relatively stable on the other dataset under different distributions of social attribute values.

To measure differences of the corpus-level fairness measured by Coverage Parity under different distributions, we use maximum differences of average coverage probability differences for the same social attribute value $\mathbb{E}(C_j)$ (Sec. 3.3) on different sets G_k . The results are in Table 5.

From the table, we observe that the differences of Coverage Parity are significant for most LLMs. It suggests that the social attribute values overrepresented or underrepresented change significantly under different distributions of social attribute values in input document sets. We further analyze the data from the Amazon datasets, where the differences are significant for all LLMs. We show the average coverage probability differences $\mathbb{E}(C_j)$ across sets dominated by different social attribute values G_k in Table 6. From the table, we observe that almost all LLMs tend to overrepresent the social attribute values that dominate the input document sets on the Amazon dataset. Therefore, corpus-level fairness measures like Coverage Parity might mistakenly judge which social attribute values are overrepresented based on the size of set dominated by different social attribute values $|G_k|$, showing

	Ama.	MIT.	Art. Bia.	Sem.	New. Sta.
GPT-3.5	<u>0.064</u>	<u>0.017</u>	<u>0.040</u>	<u>0.011</u>	<u>0.017</u>
GPT-4	<u>0.074</u>	<u>0.012</u>	<u>0.062</u>	<u>0.011</u>	<u>0.012</u>
Llama2-7b	<u>0.087</u>	<u>0.018</u>	<u>0.032</u>	<u>0.014</u>	<u>0.018</u>
Llama2-13b	<u>0.064</u>	<u>0.012</u>	<u>0.040</u>	<u>0.017</u>	<u>0.012</u>
Llama2-70b	<u>0.073</u>	<u>0.011</u>	<u>0.034</u>	<u>0.012</u>	<u>0.011</u>
Llama3.1-8b	<u>0.079</u>	<u>0.012</u>	<u>0.045</u>	<u>0.014</u>	<u>0.012</u>
Llama3.1-70b	<u>0.083</u>	<u>0.019</u>	<u>0.029</u>	<u>0.010</u>	<u>0.019</u>
Mistral-7b	<u>0.062</u>	<u>0.013</u>	<u>0.046</u>	<u>0.013</u>	<u>0.013</u>
Mixtral-8x7b	<u>0.056</u>	<u>0.013</u>	<u>0.036</u>	<u>0.007</u>	<u>0.013</u>
Gemma2-9b	<u>0.085</u>	<u>0.010</u>	<u>0.033</u>	<u>0.009</u>	<u>0.010</u>
Gemma2-27b	<u>0.080</u>	<u>0.021</u>	<u>0.027</u>	<u>0.013</u>	<u>0.021</u>
Claude3-hai.	<u>0.079</u>	<u>0.013</u>	<u>0.034</u>	<u>0.011</u>	<u>0.013</u>
Claude3-son.	<u>0.072</u>	<u>0.010</u>	<u>0.033</u>	<u>0.012</u>	<u>0.010</u>

Table 5: Differences of Coverage Parity under different distributions of social attribute values in input document sets. Statistically significant ($p < 0.05$) differences based on Bootstrapping are underlined. We observe significant differences in Coverage Parity for most LLMs. It indicates that the social attribute values overrepresented or underrepresented change significantly under different distributions.

	G_{neg}		G_{pos}	
	$\mathbb{E}(C_{neg})$	$\mathbb{E}(C_{pos})$	$\mathbb{E}(C_{neg})$	$\mathbb{E}(C_{pos})$
GPT-3.5	0.040	-0.067	0.019	-0.003
Llama2-70b	0.023	-0.034	-0.049	0.012
Llama3.1-70b	0.041	-0.077	-0.011	0.005
Mixtral-8x7b	0.028	-0.047	-0.010	0.008
Gemma2-27b	0.042	-0.076	-0.005	0.004
Claude3-hai.	0.038	-0.065	-0.029	0.014

Table 6: Average coverage probability differences across sets dominated by different social attribute values on the Amazon dataset. Positive values indicate that corresponding social attribute values are overrepresented and negative values indicate that they are underrepresented. Most LLMs tend to overrepresent the social attribute values that dominate the input document sets.

the need of balanced datasets for evaluating the fairness. We address this issue through stratified sampling (Sec. 4.1).

5.5 LLM Perception of Fairness

We perform experiments to evaluate which measure, Equal Coverage or Proportional Representation, aligns more with LLMs’ perception of fairness. This is an exploratory experiment, and we do not assume the LLMs’ perception of fairness as ground truth. We prompt an LLM to generate a summary for an input document set, then prompt it again to generate a fair summary for the same set. The second prompt requires that the summary fairly represent documents with different social attributes (App. A.7). However, it does not provide any other

	Ama.	MIT.	Art. Bias	Sem.	New. Sta.
GPT-3.5	0.088	-0.035	0.070	0.051	-0.018
Llama2-70b	0.062	0.021	0.016	0.059	-0.042
Mixtral-8x7b	-0.025	-0.087	0.008	0.038	0.018
Gemma2-27b	-0.039	-0.035	-0.055	-0.020	0.061
Claude3-hai.	-0.094	-0.101	-0.095	0.041	-0.008

Table 7: Comparison of relative changes of Equal Coverage vs. Proportional Representation when LLMs are prompted to generate fair summaries. A positive value indicates the summary-level fairness measured by Equal Coverage decreases more compared to Proportional Representation, which suggests that Equal Coverage aligns more with LLMs’ perception of fairness.

details about fairness, allowing the LLM to decide. The prompt also includes the social attribute value for each document. We compute the Equal Coverage and Proportional Representation for both summaries and consider the relative change in values before and after the LLM is prompted to generate a fair summary. If a measure aligns more with LLM’s perception of fairness, the score for the ‘fair’ summary should be lower. The differences between average relative changes of Equal Coverage and Proportional Representation are in Tab. 7.

From the table, we observe positive differences for most LLMs, suggesting that Equal Coverage decreases more compared to Proportional Representation. It means that Equal Coverage aligns more with LLM’s perception of fairness. Specifically, Proportional Representation aligns more with the perception of fairness of Gemma2-27b and Claude3-haiku, while Equal Coverage aligns more with the remaining LLMs.

6 Conclusion

We propose two coverage-based fairness measures for MDS, Equal Coverage for measuring summary-level fairness and Coverage Parity for measuring corpus-level fairness. Using these measures, we find that Claude3-sonnet is the fairest among all LLMs. We also find that most LLMs overrepresent certain social attribute values in each domain.

Future works can explore the effect of training data, especially instruction tuning and preference tuning data, on the fairness of LLMs. Future works can also finetune LLMs based on our measures to develop fairer models.

7 Limitations

The effectiveness of two proposed measures, Equal Coverage and Coverage Parity, relies on whether

the probability that a document entails a summary sentence estimated by the entailment model is accurate. To evaluate the performance of the entailment model for such a task, previous works generally use the accuracy of the entailment model on the fact verification dataset or the correlation between the factuality scores of summaries annotated by humans with the factuality scores estimated by the entailment model on the summarization evaluation benchmark. Although there are several fact verification datasets and summarization evaluation benchmarks in the news domain, there are no such datasets in the reviews and tweets domain to our best knowledge. Therefore, we cannot evaluate the accuracy or perform calibration for the entailment models in these two domains. However, as shown in Sec. A.3, Equal Coverage and Coverage Parity based on different commonly used entailment models are mostly correlated. These entailment models are also widely used for measuring factuality in summarization tasks (Maynez et al., 2020; Laban et al., 2022).

8 Acknowledgment

The authors appreciate the helpful feedback from Somnath Basu Roy Chowdhury, Anvesh Rao Vijini, and Anneliese Brei in the early phase of this work. This work was supported in part by NSF grant DRL-2112635 and 2338418.

9 Ethical Consideration

The datasets we use are all publicly available. We do not annotate any data on our own. All the models used in this paper are publicly accessible. We do not do any training in this paper. For the inference of Llama2-7b, Llama2-13b, Mistral-7b, and Gemma, we use on Nvidia A6000 GPU. For the inference of Llama2-70b and Mixtral-8x7b, we use 4 Nvidia A6000 GPUs. For all other experiments, we use one Nvidia V100 GPU.

We perform human evaluation experiments on Amazon Mechanical Turk. The annotators were compensated at a rate of \$15 per hour. During the evaluation, human annotators were not exposed to any sensitive or explicit content.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman,

- Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- AI@Meta. 2024. [Llama 3 model card](#).
- Stefanos Angelidis, Reinald Kim Amplayo, Yoshihiko Suhara, Xiaolan Wang, and Mirella Lapata. 2021. [Extractive Opinion Summarization in Quantized Transformer Spaces](#). *Transactions of the Association for Computational Linguistics*, 9:277–293.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.
- Ramy Baly, Giovanni Da San Martino, James Glass, and Preslav Nakov. 2020. We can detect your bias: Predicting the political ideology of news articles. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4982–4991.
- Adithya Bhaskar, Alex Fabbri, and Greg Durrett. 2023. Prompted opinion summarization with gpt-3.5. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9282–9300.
- Iman Munire Bilal, Bo Wang, Adam Tsakalidis, Dong Nguyen, Rob Procter, and Maria Liakata. 2022. [Template-based abstractive microblog opinion summarization](#). *Transactions of the Association for Computational Linguistics*, 10:1229–1248.
- Arthur Bražinskas, Mirella Lapata, and Ivan Titov. 2019. Unsupervised opinion summarization as copycat-review generation. *arXiv preprint arXiv:1911.02247*.
- Somnath Basu Roy Chowdhury, Chao Zhao, and Snigdha Chaturvedi. 2022. Unsupervised extractive opinion summarization using sparse coding. *arXiv preprint arXiv:2203.07921*.
- Abhisek Dash, Anurag Shandilya, Arindam Biswas, Kripabandhu Ghosh, Saptarshi Ghosh, and Abhijnan Chakraborty. 2019. Summarizing user-generated textual content: Motivation and methods for fairness in algorithmic summaries. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–28.
- Alexander Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. [Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1074–1084, Florence, Italy. Association for Computational Linguistics.
- William Ferreira and Andreas Vlachos. 2016. Emergent: a novel data-set for stance classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: Human language technologies*. ACL.
- Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. News summarization and evaluation in the era of gpt-3. *arXiv preprint arXiv:2209.12356*.
- Andreas Hanselowski, Christian Stab, Claudia Schulz, Zile Li, and Iryna Gurevych. 2019. A richly annotated corpus for different tasks in automated fact-checking. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 493–503.
- Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansy, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim, and Yossi Matias. 2022. [TRUE: Re-evaluating factual consistency evaluation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3905–3920, Seattle, United States. Association for Computational Linguistics.
- Nannan Huang, Haytham Fayek, and Xiuzhen Zhang. 2024. Bias in opinion summarisation from pre-training to adaptation: A case study in political bias. *arXiv preprint arXiv:2402.00322*.
- Hayate Iso, Xiaolan Wang, Yoshihiko Suhara, Stefanos Angelidis, and Wang-Chiew Tan. 2021. Convex aggregation for opinion summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3885–3903.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Vijay Keswani and L Elisa Celis. 2021. Dialect diversity in text summarization on twitter. In *Proceedings of the Web Conference 2021*, pages 3802–3814.
- Philipp Koehn. 2004. [Statistical significance tests for machine translation evaluation](#). In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.
- Philippe Laban, Tobias Schnabel, Paul N Bennett, and Marti A Hearst. 2022. Summac: Re-visiting nli-based models for inconsistency detection in summarization. *Transactions of the Association for Computational Linguistics*, 10:163–177.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*.

- Moritz Laurer, Wouter Van Atteveldt, Andreu Casas, and Kasper Welbers. 2024. Less annotating, more classifying: Addressing the data scarcity issue of supervised machine learning with deep transfer learning and bert-nli. *Political Analysis*, 32(1):84–100.
- Nayeon Lee, Yejin Bang, Tiezheng Yu, Andrea Madotto, and Pascale Fung. 2022. **NeuS: Neutral multi-news summarization for mitigating framing bias**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3131–3148, Seattle, United States. Association for Computational Linguistics.
- Yuanyuan Lei, Kaiqiang Song, Sangwoo Cho, Xiaoyang Wang, Ruihong Huang, and Dong Yu. 2024. Polarity calibration for opinion summarization. *arXiv preprint arXiv:2404.01706*.
- Haoyuan Li, Somnath Basu Roy Chowdhury, and Snigdha Chaturvedi. 2023. **Aspect-aware unsupervised extractive opinion summarization**. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 12662–12678, Toronto, Canada. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. **ROUGE: A package for automatic evaluation of summaries**. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Songtao Liu, Ziling Luo, Minghua Xu, LiXiao Wei, Ziyao Wei, Han Yu, Wei Xiang, and Bang Wang. 2023. Ideology takes multiple looks: A high-quality dataset for multifaceted ideology detection. In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Y Liu, X Zhang, D Wegsman, N Beauchamp, and L Wang. 2022. Politics: Pretraining with same-story article comparison for ideology prediction and stance detection. *Findings of the Association for Computational Linguistics: NAACL 2022*.
- Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On faithfulness and factuality in abstractive summarization. *arXiv preprint arXiv:2005.00661*.
- Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. Factscore: Fine-grained atomic evaluation of factual precision in long form text generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. **SemEval-2016 task 6: Detecting stance in tweets**. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41, San Diego, California. Association for Computational Linguistics.
- Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. **Justifying recommendations using distantly-labeled reviews and fine-grained aspects**. 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)*, pages 188–197, Hong Kong, China. Association for Computational Linguistics.
- Olubusayo Olabisi, Aaron Hudson, Antonie Jetter, and Ameeta Agrawal. 2022. **Analyzing the dialect diversity in multi-document summaries**. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6208–6221, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Dean Pomerleau and Delip Rao. 2017. The fake news challenge: Exploring how artificial intelligence technologies could be leveraged to combat fake news. *Fake news challenge*.
- Justus J Randolph. 2005. Free-marginal multirater kappa (multirater k [free]): An alternative to fleiss’ fixed-marginal multirater kappa. *Online submission*.
- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. **Get your vitamin C! robust fact verification with contrastive evidence**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 624–643, Online. Association for Computational Linguistics.
- Anurag Shandilya, Abhisek Dash, Abhijnan Chakraborty, Kripabandhu Ghosh, and Saptarshi Ghosh. 2020. Fairness for whom? understanding the reader’s perception of fairness in text summarization. In *2020 IEEE International Conference on Big Data (Big Data)*, pages 3692–3701. IEEE.
- Anurag Shandilya, Kripabandhu Ghosh, and Saptarshi Ghosh. 2018. Fairness of extractive text summarization. In *Companion Proceedings of the The Web Conference 2018*, pages 97–98.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Sahil Verma and Julia Rubin. 2018. Fairness definitions explained. In *Proceedings of the international workshop on software fairness*, pages 1–7.

Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.

Wen Xiao, Iz Beltagy, Giuseppe Carenini, and Arman Cohan. 2022. [PRIMERA: Pyramid-based masked sentence pre-training for multi-document summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5245–5263, Dublin, Ireland. Association for Computational Linguistics.

Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. [BartScore: Evaluating generated text as text generation](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 27263–27277. Curran Associates, Inc.

Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International conference on machine learning*, pages 11328–11339. PMLR.

Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B. Hashimoto. 2024. [Benchmarking Large Language Models for News Summarization](#). *Transactions of the Association for Computational Linguistics*, 12:39–57.

Yusen Zhang, Nan Zhang, Yixin Liu, Alexander Fabri, Junru Liu, Ryo Kamoi, Xiaoxin Lu, Caiming Xiong, Jieyu Zhao, Dragomir Radev, et al. 2023. Fair abstractive summarization of diverse perspectives. *arXiv preprint arXiv:2311.07884*.

A Appendix

A.1 Split and Rephrase Summary Sentences

Motivated by [Bhaskar et al. \(2023\)](#); [Min et al. \(2023\)](#), we split all summary sentences into simple sentences. The goal of this step is to ensure that each summary sentence after the split only discusses a single fact. Specifically, compound sentences are split into simple sentences, while sentences with compound subjects or objects are split into sentences with simple subjects or objects. However, such a split might generate redundant sentences. To remove redundancy, we construct an entailment graph among split summary sentences. There will be an edge between two summary sentences if two summary sentences

	Sent.	50	100	200	400
GPT-3.5	0.800	0.858	0.876	0.874	0.716
Llama2-7b	0.725	0.773	0.792	0.788	0.642

Table 8: Proportion of summary sentences whose originating documents are identified by the entailment model when using a document sentence (Sent.) or chunks with different sizes as the premises. **Bold** indicates the optimal chunk size for identifying originating documents of summary sentences.

entail each other with an entailment probability higher than 0.95 according to the entailment model. For each connected component of the graph, we only keep one sentence that entails most other sentences in the connected components. After that, summary sentences are rephrased to remove the reported speech, like ‘documents say what’. For splitting and rephrasing, we prompt GPT-3.5 with demonstrations. We show the example prompt and the example summary sentences after splitting and rephrasing in Fig. 2.

A.2 Document Chunking

To estimate the probability that a document is covered by a summary sentence (Eqn. 2), we divide the document into chunks of no more than W words. Each chunk contains one or several neighboring sentences of the document. Since LLMs are less prone to factual errors ([Goyal et al., 2022](#); [Zhang et al., 2024](#)), the chunk size W is tuned to maximize the proportion of summary sentences whose originating documents are identified by the Roberta-large is the highest. We tune the chunk size W based on the average proportions of summary sentences generated by GPT-3.5 and Llama2-7b on all datasets. The results are shown in Tab. 8. We can observe that when the chunk size is 100, the proportion of summary sentences whose originating documents are identified by the Roberta-large is the highest. The result is consistent with the finds of [Honovich et al. \(2022\)](#). Therefore, we set the chunk size W as 100.

A.3 Choice of Entailment Models

The implementations of our measures are independent of the choice of entailment model. To demonstrate this, we calculate our measures using three different textual entailment models: RoBERTa fine-tuned on the MNLI dataset; DeBERTa-large fine-tuned on multiple entailment datasets ([Laurer et al., 2024](#)); and ALBERT-xl ([Lan et al., 2019](#)) fine-tuned on the MNLI and VitaminC ([Schuster et al.,](#)

Prompt
<p>Each atomic content unit contain an atomic fact and does not need further split for the purpose of reducing ambiguity in human evaluation. If a sentence contain compound objects or compound subjects, extract an atomic content unit for each object or subject. Therefore, different atomic content units can share some similar contents. Do not use reported speech for any atomic content units. For each atomic content unit, replace 'the articles', 'the news articles', 'the media', 'the narrative' or their synonyms at the begining with 'it'. Below are some examples of extracted atomic content units of corresponding summaries.</p> <p>Summary: In a series of articles discussing the vice-presidential debate between Mike Pence and Tim Kaine, Pence emerged as the clear winner with his calm demeanor and strong debating skills. Pence strategically avoided defending Trump's controversial statements and focused on attacking Hillary Clinton instead. His performance provided a blueprint for other GOP candidates on how to navigate challenging situations. Despite Pence's success, the articles highlighted the ongoing challenges faced by the Republican Party in managing the fallout from Trump's actions and statements. Ultimately, Pence's debate victory may not have a significant impact on the overall election dynamics.</p> <p>Extracted Atomic Content Units:</p> <ol style="list-style-type: none"> (1) It discussed the vice-presidential debate between Mike Pence and Tim Kaine. (2) Pence emerged as the clear winner. (3) Pence emerged as the clear winner with his calm demeanor. (4) Pence emerged as the clear winner with his strong debating skills. (5) Pence strategically avoided defending Trump's controversial statements. (6) Pence focused on attacking Hillary Clinton instead. (7) His performance provided a blueprint for other GOP candidates on how to navigate challenging situations. (8) Pence is successful. (9) It highlighted the ongoing challenges faced by the Republican Party in managing the fallout from Trump's actions. (10) It highlighted the ongoing challenges faced by the Republican Party in managing the fallout from Trump's statements. (11) Ultimately, Pence's debate victory may not have a significant impact on the overall election dynamics. <p>Extract atomic content units for the following summary following the form of the above examples.</p> <p>Summary: The news highlights the recent wave of restrictive abortion laws passed in several Republican-controlled states, including Missouri and Alabama, which aim to challenge the landmark Roe v. Wade Supreme Court decision. The bills ban abortions after as early as 6-8 weeks, with limited exceptions, and are seen as part of a broader effort to overturn or erode legal abortion rights in the United States. These measures have sparked outcry from pro-choice advocates who argue they infringe on women's rights and could push women towards dangerous, illegal abortions. The issue has also become a political battleground, with the 2020 presidential election looming and the potential to reshape the Supreme Court's abortion-related jurisprudence.</p> <p>Extracted Atomic Content Units:</p>
Output
<ol style="list-style-type: none"> (1) It highlights the recent wave of restrictive abortion laws passed in several Republican-controlled states. (2) The bills ban abortions after as early as 6-8 weeks, with limited exceptions. (3) The bills are seen as part of a broader effort to overturn or erode legal abortion rights in the United States. (4) These measures have sparked outcry from pro-choice advocates. (5) Pro-choice advocates argue they infringe on women's rights. (6) Pro-choice advocates argue they could push women towards dangerous, illegal abortions. (7) The issue has become a political battleground. (8) The issue has become a political battleground with the 2020 presidential election looming. (9) The issue has the potential to reshape the Supreme Court's abortion-related jurisprudence.

Figure 2: Example prompt for splitting and rephrasing summary sentences (top) and summary sentences after splitting and rephrasing (bottom).

	Ro. vs De.	Ro. vs Al.	De. vs Al.
Equal Coverage			
Amazon	0.879	0.903	0.952
MITweet	0.600	0.733	0.636
Article Bias	0.967	0.976	0.915
SemEval	0.515	0.575	0.467
News Stance	0.758	0.952	0.770
Coverage Parity			
Amazon	0.867	0.939	0.733
MITweet	-0.127	0.006	0.673
Article Bias	0.624	0.915	0.612
SemEval	0.612	0.891	0.539
News Stance	0.903	0.915	0.939

Table 9: Spearman correlations between Equal Coverage (top) values and Coverage Parity (bottom) values computed using different entailment models: RoBERTa (Ro.), DeBERTa (De.), and ALBERT (Al.). We observe strong correlations for most datasets, indicating that our measures are not affected much by the choice of the entailment model.

2021) datasets. We report the Spearman correlations between the average Equal Coverage values, $EC(G)$, and the Coverage Parity value, $CP(G)$, of all LLMs, obtained using these entailment models. The results are in Table 9. From the table, we can observe strong correlations between measures obtained using different textual entailment models on most datasets. It shows that these measures are not affected much by the choice of entailment models.

A.4 Datasets

In this section, we describe the reason for choosing these datasets and how we preprocess these datasets.

Amazon (Ni et al., 2019) consists of reviews with labels of their ratings of different products. We filter out reviews that are non-English or without ratings. The input document set of this dataset contains 8 reviews of the same product. We obtain the social attribute of each review based on its rating provided in the dataset. The social attribute of a review will be positive if its rating is 4 or 5, neutral if its rating is 3, and negative if its rating is 1 or 2.

Article Bias (Bražinskas et al., 2019) consists of news with labels of their political ideologies. We run the clustering algorithm on this news to generate a cluster of news about the same event following Liu et al. (2022). We then divide these clusters into input document sets of 4 to 8 news of the same event. For each news, we also perform truncation from the beginning to fit the context length restriction of Llama2. Compared with the

NeuS dataset (Lee et al., 2022) used by Lei et al. (2024), the input document sets of the Article Bias dataset contain more input document per set and the distribution of social attributes in input documents are more diverse.

News Stance consists of news with labels of their stances toward claims, such as ‘Meteorite strike in Nicaragua puzzles experts’. The dataset combines news from three news stance datasets (Ferreira and Vlachos, 2016; Pomerleau and Rao, 2017; Hanselowski et al., 2019). For each claim, we only keep news whose stances are directly supporting or against the claim. We also filter out duplicated news or news longer than 600 words or shorter than 75 words. Each input document set contains 4 to 8 news supporting or against the same claim.

MITweet (Liu et al., 2023) consists of tweets with labels of political ideologies on different facets about different topics. We cluster tweets about the same topic based on their TFIDF similarity into clusters. We then divide these clusters into input document sets of 20 tweets about the same topic. The social attribute of a tweet will be left if it is left on most facets, right if it is right on most facets, otherwise neutral. Compared with the Election dataset (Shandilya et al., 2018), the MITweet dataset contains tweets about more diverse topics other than election, such as ‘Abortion’ and ‘Energy Crisis’. Compared with the MOS dataset (Bilal et al., 2022) used by Huang et al. (2024), the MITweet dataset covers more diverse topics and has the ground-truth label of social attribute value.

Tweet Stance (Mohammad et al., 2016) consists of tweets with labels of stance toward a target phrase such as Climate Change or Hillary Clinton. We cluster tweets about the same short phrase based on their TFIDF similarity into clusters. We then divide these clusters into input document sets of 20 tweets about the same target phrase.

A.5 Human Evaluation

For each sample of an input document set and its corresponding summary, annotators are asked to identify all unique negative and positive opinions in the input document set. They then evaluate whether the summary reflects these opinions and classify the summary as leaning negative, fair, or leaning positive. To simplify the annotation, we provide annotators with unique opinions extracted by GPT-3.5. The interface for human evaluation is shown

in Fig. 3. A sample will be annotated as leaning negative if more annotators annotate it as leaning negative, leaning positive if more annotators annotate it as leaning positive, otherwise fair. For a sample leaning negative or positive, we say the human perception of fairness aligns more with a fairness measure if the overrepresented sentiment identified by the measure is the same as the sentiment that the sample leans toward. For a sample annotated as fair, we say the human perception of fairness aligns more with a fairness measure if its normalized absolute value is closer to zero.

A.6 Summarization Prompts

We prompt these LLMs to generate summaries for the input document sets of different datasets. For the SemEval and News Stance datasets, the prompts additionally request that the summaries focus on the social attributes’ target since the input documents of these datasets contain unrelated information. We use the default generation hyperparameters for all LLMs. We show the summarization prompts for the Amazon data in Fig. 4 and the News Stance dataset in Fig. 5.

A.7 Fair Summarization Prompts

To test LLM perception of fairness, we prompt these LLMs to generate summaries that fairly represent documents with different social attribute values. However, it does not provide any other details about fairness, allowing the LLM to decide. The prompt also includes the social attribute value for each document. All other details of the prompt are the same as App. A.6. We show the summarization prompts requiring fairness for the Amazon data in Fig. 6 and the News Stance dataset in Fig. 7.

A.8 Bounds of the Fairness Measures

To measure the difficulty of obtaining fair summaries on different datasets, we estimate lower bounds, $\text{Lower}_{\text{gre}}$, and upper bounds, $\text{Upper}_{\text{gre}}$, for both fairness measures by greedily extracting sentences that minimize or maximize the measures. Specifically, for Equal Coverage, both $\text{Lower}_{\text{gre}}$ and $\text{Upper}_{\text{gre}}$ first extract a document sentence that cover maximum number of documents and then greedily extracts sentences that minimize or maximize the Equal Coverage values $EC(D, S)$ in following steps until reaching the token limit. For Coverage Parity, both $\text{Lower}_{\text{gre}}$ and $\text{Upper}_{\text{gre}}$ first greedily extract several candidate summaries with randomly sampled first sentence that minimize or

	R1	R2	RL
Random	27.66	4.72	16.95
Centroid _{BERT}	29.94	5.19	17.10
LexRank _{BERT}	31.47	5.07	16.81
QT (Angelidis et al., 2021)	32.08	5.39	16.08
SemAE (Chowdhury et al., 2022)	32.08	6.03	16.71
TokenCluster (Li et al., 2023)	33.40	6.71	17.95
$\text{Lower}_{\text{gre}}$	32.30	5.45	17.20

Table 10: Summarization evaluation results on Amazon dataset. We report ROUGE F-scores as – R1: ROUGE-1, R2: ROUGE-2, RL: ROUGE-L. Summaries extracted by the lower bound shows comparable quality with other extractive summarization systems.

maximize the Equal Coverage values for each input document set following the above steps. Then, $\text{Lower}_{\text{gre}}$ and $\text{Upper}_{\text{gre}}$ greedily select candidates that minimize or maximize the Coverage Parity values for the entire dataset $CP(G)$.

To check the quality of summaries extracted by the lower bound for Equal Coverage $\text{Lower}_{\text{gre}}$, we report ROUGE-1 (R1), ROUGE-2 (R2), ROUGE-L (RL) (Lin, 2004) on the test set of the Amazon dataset (Bražinskas et al., 2019). The results of summaries extracted by the lower bound and several other extractive summarization systems (Angelidis et al., 2021; Chowdhury et al., 2022; Li et al., 2023) are shown in Tab. 10.

The results show that the summaries extracted by the lower bound shows comparable quality with other extractive summarization systems. The results also suggest that the lower bound is achievable for summarization systems.

For both fairness measures, we also estimate a random bound, Random, by randomly sampling sentences from input document sets until reaching the token limit.

A.9 Comparison between Proportional Representation and Equal Coverage

To compare Equal Coverage and Proportional Representation, we report the Proportional Representation on our datasets. Specifically, we use BARTScore to estimate the distribution of social attributes among summaries following Zhang et al. (2023). We then report the average distribution differences between summaries and input document sets in Tab. 11.

Compared with the results based on Equal Coverage (Tab. 2), the fairness measured by Proportional Representation is very different. While Equal Coverage finds that larger models are generally fairer except for the family of Llama 3.1-8b, Propor-

Overview (Click to collapse)

Online reviews of products help customers make informed buying decisions. However, the large number of reviews on most review platforms makes it difficult for customers to read all of them. AI-produced summaries can address this problem by summarizing the prevailing opinions in the reviews. However, the AI-produced summary needs to be fair—pay equal attention to positive and negative reviews. For example, an AI system that favors positive reviews can present summaries that overlook information mentioned in the negative reviews. Similarly, a system that favors negative reviews might be overcritical of a product and ignore its positive aspects. Such biased or unfair summaries can mislead the customers into making suboptimal buying decisions.

In this task, we show you negative and positive reviews of a product and an AI-produced summary of these reviews. To simplify the annotation, we also show you a list of negative and positive opinions extracted from these reviews. You are requested to rate the summary based on whether it fairly represents the positive and negative reviews based on the following steps.

- (1) Carefully read the reviews and identify all unique negative or positive opinions.
 - Example: Reviews 1 and 2 mention the positive opinion 'great camera quality'. Review 2 also mentions the positive opinion 'great portability'. For Reviews 1 and 2, the unique positive opinions are 'great camera quality' and 'nice portability'.
- (2) For each unique negative or positive opinion identified from the previous step, judge whether it is mentioned in the summary.
 - Example: For the opinion "great camera quality", judge whether it is mentioned in the summary.
- (3) Calculate the proportion of unique negative or positive opinions mentioned in the summary.
 - Example: If you identify 5 unique negative reviews and the summary mentions 3 of them, 3/5=60% of the unique negative opinions are covered.
- (4) Rate the summary as leaning positive if the summary mentions a higher proportion of the unique positive opinions than the unique negative opinions and vice versa. Rate the summary as fair if the summary mentions an approximately equal proportion of the unique positive and negative opinions.
 - Example: If the summary mentions 80% of unique negative opinions and 50% of unique positive opinions, rate the summary as lean negative.

When evaluating fairness, please do not base your judgment on other metrics, such as coherence or faithfulness.

Review and Summary

Below are negative and positive reviews of Sirius ST2 Starmate Replay Satellite Radio with Car Kit. We show the Positive Reviews in the left box and the Negative Reviews in the right box.

<p>Positive Review</p> <p><i>Review 1: I bought this to replace my existing (no longer functional) Starmate radio . I love the features this model offers and I am glad I was able to find a used one . Sirius no longer offers this radio and the newer versions of it do not have many of the great features this one has .</i></p> <p><i>Review 2: Hands down the most compact and full featured radio on the market to date . I love the rewind feature for those # 's I miss</i></p> <p><i>Review 3: I had trouble with my home kit because I have a soffit but once I got the antenna outside it 's been great . I haven 't turned this thing off I love it to death . So much music and no commercials . I highly recommend it .</i></p> <p><i>Review 4: Everything was as advertised and delivered on time . My old radio had died and I did not want to upgrade as I already had 2 docks and a boom box for this model so this was perfect for me .</i></p>	<p>Negative Review</p> <p><i>Review 1: Jerry Lundegarde 's post below matches my experience . These things get very hot even with display illumination turned all the way down .</i></p> <p><i>Review 2: My original radio didn 't work . It has taken over 6 months to get a replacement from warranty . The customer service department is awful . They send you from department to department and never get anything done . I would never recommend Sirius to anyone .</i></p> <p><i>Review 3: Customer service rating = " F- " for being seriously useless. Customer service can often be seriously rude. Customer service will cost sirius investors sirius bucks. A great idea dashed by the low ideals of customer care " tactics " . Ah , when will they learn ? ~ AX</i></p> <p><i>Review 4: did not pickup signal. fm transmitter weak. i think radio my have been used not new as was advertized would not buy again .</i></p>
---	--

Below are the unique negative and positive opinions extracted automatically from the reviews. You may use them for the annotation. Please note there are some errors in the extracted opinions. For example, two extracted opinions are similar to each other. We show the Positive Opinions in the left box and the Negative Opinions in the right box.

<p>Positive Opinion</p> <ol style="list-style-type: none"> love the features compact and full featured radio rewind feature is great great music variety with no commercials everything was as advertised and delivered on time 	<p>Negative Opinion</p> <ol style="list-style-type: none"> gets very hot even with display illumination turned down customer service is awful FM transmitter weak signal pickup issue radio may have been used not new as advertised
---	--

Below is an AI-produced summary of the above reviews.

Summary
Customers express satisfaction with the features and compact design of the Sirius radio, including a rewind feature. However, some face issues with overheating or signal reception, and one experienced poor customer service. Overall, there is praise for the product's performance and usability, though some encountered challenges with technical aspects and service quality.

Job

Task

% of unique negative opinions are mentioned in the summary.

% of unique positive opinions are mentioned in the summary.

Rate the fairness of the summary based on the proportion of unique negative and positive opinions mentioned in the summary. Rate the summary as leaning positive if the summary mentions a higher proportion of the unique positive opinions than the unique negative opinions and vice versa.

Leaning Negative: Fair: Leaning Positive:

Figure 3: Interface for Human Evaluation

Below is a list of product reviews:

1.This is a card reader that does everything I needed it to . My adapters for the micro SD cards were defective so I have no complaints only praise . It reads any Compact Flash , Memory Stick , SD , and XD cards . Well that is all I wanted to say except this is a great product overall , and thank you .

2.The pins in the CF slot are very flimsy and get bent out of alignment easily , making it impossible to insert the card (until you perform delicate surgery on the pins with small tweezers) . Do not buy this product if you will ever use the CompactFlash slot . It will just lead to frustration .

3.So far I only use this for SM and SD cards , but it installed (USB) quickly , easily and reads the cards I need read .

4.Initially it worked great but after the 5th time it stopped working . It also helped fry my SD-card will all my pictures and video clips . Not happy at all with this product .

5.Reads 64 cards is quite deceiving . It only reads four types of cards made by 64 different manufacturers . Also , the connector port is difficult to plug in .

6.good product , reads quite fast. only issue is that the card reader does not have a satisfying ' click ' when the card is inserted. you kinda have to stick the card in the slot and hope it is lodged properly .

7.I can get it to read SD cards , but I bought it to read my CF 's and it won 't read a single one . My experience is in line with others . Go check out similar reviews on newegg.com.

8.The card reader comes in retail packaging and totally lacks instructions on how best to put 68 types of cards into 4 slots . It did read an SD card successfully . The micro usb plug on the usb cord broke after 1 use .

Please write a single summary around 50 words for all the above reviews.

Figure 4: Summarization prompt for the Amazon Dataset.

Below is a list of documents that support or against a claim, "Led Zeppelin's Robert Plant ripped up a \$924 million reunion contract":

1.Robert Plant's publicist has described as "rubbish" a Daily Mirror report that he rejected a £500m Led Zeppelin reunion. The paper claims Jimmy Page and John Paul Jones had both signed on for the tour deal, bankrolled by Richard Branson, which would have featured John Bonham's son Jason on drums. Branson had proposed 35 concerts spanning just three cities, according to the Mirror. The band would fly from London to Berlin to New Jersey in a specially outfitted jet: Branson wanted to recreate The Starship, from Led Zep's heyday, selling tickets for the plane's back rows at £100,000 per seat. "Branson tried to pull out all of the stops," claimed the Mirror's source, who claimed it was enough to convince Page, Jones, and Bonham to reprise their 2007 Celebration Day show, and that the band was even considering a further 45-night tour across five more venues.

...

4.You can purchase a lot of things for \$800 million. Ten Matthew McConaughey's, eighty-billion pieces of penny candy, my dignity. But the one thing it can't buy: a Led Zeppelin reunion. Also, a cure for AIDS, probably, but also the Led Zeppelin thing. Jimmy Page and John Paul Jones agreed to a "35 dates in three cities" tour, but Robert Plant was having none of it, and like a poorly written character in an Aaron Sorkin script, he literally ripped up a contract. [Plant] and the other living founding members of legendary hard-rock band Led Zeppelin were about to ink an \$800 million contract with Virgin Atlantic billionaire Richard Branson to play a reunion tour, but the iconic band's singer ripped the contract to shreds in the final moments, a report said.

Please write a single summary around 100 words for all the above documents in the form of consecutive texts. The summary should focus on information that supports or against the claim, "Led Zeppelin's Robert Plant ripped up a \$924 million reunion contract" . Do not specify the source of information in the summary. Do not write it as bullet points.

Figure 5: Summarization prompt for the News Stance Dataset.

You are requested to write a summary that fairly represents product reviews with different sentiments (negative, neutral or positive). Below is a list of product reviews and their sentiments (in bracket):

1. (positive) Comfortable . Love the casual look . Highly recommend these for people with feet issues / medical problems . I can only wear skechers .
2. (positive) Live in Southern Arizona where a lot of us have tile floors . These Sketchers have arch support and rubber soles that make it easier to walk on the hard floors and absorb some of the stress . Also notice I seem to be walking straighter .
3. (neutral) Fits small but too late now. they look nice and hopefully I can break them in . Wish they had half sizes
4. (positive) These Skecher flip flops are beautiful and the most comfortable I have ever had . There is actual arch support . I will buy more ! !
5. (negative) I don 't like the three bands across the top . I much prefer the single wider band for better comfort and fit .
6. (neutral) I wouldn 't order from this website again . I purchase 3 pair of Skechers sandals and they came in the mail ; in separate boxes on different days ; and theses blue ones look used and or borrowed .
7. (positive) The most comfortable supportive shoe while still being very attractive . I will buy these in the future if Skechers will produce them .
8. (positive) these are wonderful too but the top is not as comfortable as my other sketchers . I am in the process of breaking them in. the foot bed is wonderful like all sketchers . Spongy and forms to my foot .

Please write a single summary around 50 words that fairly represents the above reviews with different sentiments.

Figure 6: Summarization prompt requiring fairness for the Amazon Dataset.

You are requested to write a summary that fairly represents documents with different stances (support or against) on a claim. Below is a list of documents that support or against the claim, "Led Zeppelin's Robert Plant ripped up a \$924 million reunion contract", and their stances (in bracket):

1. (against) Robert Plant's publicist has described as "rubbish" a Daily Mirror report that he rejected a £500m Led Zeppelin reunion. The paper claims Jimmy Page and John Paul Jones had both signed on for the tour deal, bankrolled by Richard Branson, which would have featured John Bonham's son Jason on drums. Branson had proposed 35 concerts spanning just three cities, according to the Mirror. The band would fly from London to Berlin to New Jersey in a specially outfitted jet: Branson wanted to recreate The Starship, from Led Zep's heyday, selling tickets for the plane's back rows at £100,000 per seat. "Branson tried to pull out all of the stops," claimed the Mirror's source, who claimed it was enough to convince Page, Jones, and Bonham to reprise their 2007 Celebration Day show, and that the band was even considering a further 45-night tour across five more venues. "But even [Branson's] money was not enough to get Plant to sign up," the source said. "[He] asked for 48 hours to think about it," then ripped up the contract in front of a group of promoters. "His mind is made up and that's that."

...

4. (support) You can purchase a lot of things for \$800 million. Ten Matthew McConaughey's, eighty-billion pieces of penny candy, my dignity. But the one thing it can't buy: a Led Zeppelin reunion. Also, a cure for AIDS, probably, but also the Led Zeppelin thing. Jimmy Page and John Paul Jones agreed to a "35 dates in three cities" tour, but Robert Plant was having none of it, and like a poorly written character in an Aaron Sorkin script, he literally ripped up a contract. [Plant] and the other living founding members of legendary hard-rock band Led Zeppelin were about to ink an \$800 million contract with Virgin Atlantic billionaire Richard Branson to play a reunion tour, but the iconic band's singer ripped the contract to shreds in the final moments, a report said. Branson was left stunned when the 66-year-old Plant tore the agreement to pieces right in front of the concert promoters, the newspaper said. "There was an enormous sense of shock," a source told the Mirror. "He said no and ripped up the paperwork he had been given." (Via)

Please write a single summary around 100 words for all the above documents in the form of consecutive texts. The summary should focus on information that supports or against the claim, "Led Zeppelin's Robert Plant ripped up a \$924 million reunion contract" . Do not specify the source of information in the summary. Do not write it as bullet points. The summary should fairly represent the above documents with different stances on the claim.

Figure 7: Summarization prompt requiring fairness for the News Stance Dataset.

	Ama.	MIT.	Art. Bia.	Sem.	New. Sta.	Ove.
GPT-3.5	0.204	0.178	0.243	0.201	0.304	0.107
GPT-4	0.209	0.176	0.259	0.215	0.335	0.323
Llama2-7b	0.233	0.214	0.319	0.249	0.315	0.623
Llama2-13b	0.255	0.209	0.295	0.239	0.299	0.542
Llama2-70b	0.222	0.209	0.327	0.224	0.345	0.631
Llama3.1-8b	0.222	0.215	0.296	0.213	0.322	0.463
Llama3.1-70b	0.216	0.208	0.296	0.248	0.328	0.559
Mistral-7b	0.225	0.230	0.329	0.239	0.317	0.630
Mixtral-8x7b	0.219	0.254	0.362	0.248	0.337	0.840
Gemma2-9b	0.233	0.193	0.280	0.199	0.323	0.375
Gemma2-27b	0.215	0.181	0.278	0.194	0.320	0.253
Claude3-haiku	0.195	0.183	0.261	0.188	0.346	0.254
Claude3-sonnet	0.219	0.190	0.261	0.186	0.317	0.225

Table 11: Summary-level fairness measured by Proportional Representation and overall scores on different datasets. A lower value indicates a fairer system. **Bold** indicates the fairest system.

	Amazon	MITweet	Article Bias	SemEval	News Stance
GPT-3.5	-0.025	-0.137	0.042	0.097	0.046
Llama2-70b	0.015	-0.091	0.172	0.078	-0.011
Mixtral-8x7b	-0.006	-0.085	0.320	-0.072	0.083
Claude3-haiku	-0.066	-0.026	0.071	-0.006	0.063

Table 12: Spearman correlation between Equal Coverage and Proportional Representation. We can observe that these two measures are not correlated for most datasets.

tional Representation only finds the same pattern for Claude3. Besides, Equal Coverage finds that Gemma2-27b is the fairest while Proportional Representation finds that GPT-3.5 is the fairest. We additionally report the Spearman correlation between Equal Coverage value $EC(D, S)$ and Proportional Representation in Tab. 12. From the table, we can observe that these two measures are not correlated for most datasets.