

# Summarization of Opinionated Political Documents with Varied Perspectives

Nicholas Deas, Kathleen McKeown

Department of Computer Science, Columbia University

[ndeas,kathy]@cs.columbia.edu

## Abstract

Global partisan hostility and polarization has increased, and this polarization is heightened around presidential elections. Models capable of generating accurate summaries of diverse perspectives can help reduce such polarization by exposing users to alternative perspectives. In this work, we introduce a novel dataset and task for independently summarizing each political perspective in a set of passages from opinionated news articles. For this task, we propose a framework for evaluating different dimensions of perspective summary performance. We benchmark 11 summarization models and LLMs of varying sizes and architectures through both automatic and human evaluation. While recent models like GPT-4o perform well on this task, we find that all models struggle to generate summaries that are faithful to the intended perspective. Our analysis of summaries focuses on how extraction behavior is impacted by features of the input documents<sup>1</sup>.

## 1 Introduction

Political ideologies can lead people to develop misperceptions of groups with opposing opinions (Chambers et al., 2006), and political events, such as the 2024 US presidential election, French legislative election, or the Brexit referendum, can reinforce negative attitudes (Hanna et al., 2013; Wellings et al., 2024). These misperceptions, reinforced by news consumption on social media (Levy, 2021), contribute to increased polarization and instability (Braley et al., 2023; Lees and Cikara, 2021). Exposure to alternative perspectives, however, has been shown to help alleviate polarization (Baliotti et al., 2021). To encourage such exposure, some groups focus on aggregating (e.g., All Sides<sup>2</sup>) or summarizing different political per-

<sup>1</sup>We make our code available at <https://github.com/NickDeas/Polisum>

<sup>2</sup><https://www.allsides.com/>

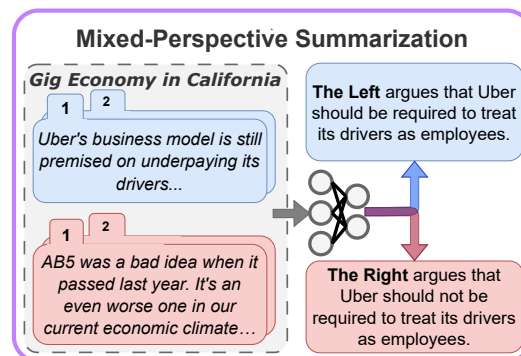


Figure 1: Example of the Mixed-Perspective Setting of POLISUM. Given a set of opinionated passages reflecting a mix of views, POLISUM requires models to generate both a summary of the left and right-leaning political perspectives.

spectives on divisive issues (e.g., The Flip Side<sup>3</sup>, Ground News<sup>4</sup>).

Large language models' (LLMs) ability to summarize opinions and news has recently neared human performance (Zhang et al., 2023a; Bhaskar et al., 2023). Recent work, however, has also shown that LLMs can unfairly represent diverse opinions in review and tweet summarization (Zhang et al., 2023b; Huang et al., 2023; Tay, 2019).

To evaluate current models' capabilities in summarizing multiple perspectives in a set of input documents, we propose a novel dataset, POLISUM, and task, mixed perspective summarization<sup>5</sup>. Our task involves generating independent abstractive summaries of both left and right-leaning political perspectives given collections of opinionated news passages with mixed perspectives. Mixed perspective summarization thus requires models to both

<sup>3</sup><https://www.TheFlipSide.io/>

<sup>4</sup><https://ground.news/>

<sup>5</sup>We make our dataset available to researchers who have been granted permission to use the data for their work by The Flip Side.

distinguish the perspectives reflected in input texts as well as produce a pair of summaries representing those differing perspectives. We summarize our primary contributions as follows:

1. We introduce a novel benchmark dataset, POLISUM, for mixed perspective summarization, requiring models to independently summarize political perspectives on controversial issues, provided passages from news and op-eds.
2. We introduce an initial framework for evaluating different components of perspective summaries, including both human evaluation and automatic metrics.
3. We benchmark 11 models, finding that most approaches struggle to faithfully summarize the intended perspective and that improving automatic evaluation metrics is an important area for future work.
4. We analyze model extraction behavior (e.g., which documents models extract from and to what degree) and find that generated summaries suffer from biases due to input document position, document length, and use of arousing terms.

## 2 POLISUM Dataset

**Data Source.** To enable evaluating political perspective summaries, we collect opinionated texts and paired summaries from The Flip Side<sup>6</sup>. Since September 2017, The Flip Side has covered politically controversial issues (i.e. overturning of *Roe v. Wade*, presidential elections) in the United States. Toward bridging political divides, each publication includes a headline reflecting the central issue, curated passages authored by political commentators with varying political leaning, and a pair of short summaries of the left and right perspectives<sup>7</sup>. We collect this information from The Flip Side archive, gathering nearly 1,000 samples. An excerpted example of source passages and summaries is shown in Table 2.

Table 1 presents summary statistics of the collected multi-document, multi-summary dataset (data statement included in Appendix C). We also

<sup>6</sup>We thank The Flip Side for generously allowing us to use their archive for this work: <https://www.TheFlipSide.io/archives>

<sup>7</sup>Some publications lack individual perspective summaries, primarily on topics where both parties largely agree or for meta-articles that do not cover political topics. These samples are not included in the dataset.

show statistics for CoCoTRIP (Iso et al. 2022; the most comparable multi-document, multi-summary dataset), XSUM (Narayan et al. 2018; a single-document, single-summary corpus with similarly short references), and *Multi-News* (Fabbri et al. 2019, a multi-document, single-summary news corpus). The aim of CoCoTRIP is to generate multiple hotel review summaries for a pair of comparable hotels: two summaries highlighting each hotel’s unique attributes, and one highlighting similarities. In contrast, POLISUM contains almost 20 times more labeled examples and focuses on political perspectives rather than reviews. Additionally, compared to XSum, POLISUM is similar in length but slightly more abstractive, measured by n-gram overlap between source and reference. On average, there are 9.6 source passages for each publication (4.8 per perspective), with a combined length comparable to a news article (~1,300 tokens).

**Perspective Summarization Task.** Using this dataset, the aim of our proposed task is to generate a pair of summaries (i.e., *Summaries* in Table 2) representing the left and right political perspectives on a controversial topic (i.e. *Headline*). As input, models are provided with two sets of opinionated editorial passages (i.e., *Source Passages*), each representing one of the two political perspectives. Thus, the dataset enables a multi-document, multi-summary task. As with traditional news summarization tasks, the summary for each perspective should condense the information present in the passages. At the same time, the generated left and right-perspective summaries should faithfully represent the perspectives reflected in their respective passages.

## 3 Experiments

Traditional opinion summarization typically involves summarizing a common perspective among all inputs, which we refer to as the Single-Perspective setting<sup>8</sup>. We instead introduce an additional challenge requiring models to identify input documents relevant to the intended perspective, referred to as the Mixed-Perspective setting and illustrated in Figure 1. In this setting, approaches are provided with a mixed set of left and right-leaning passages as input and are tasked with producing independent left and right-perspective summaries from the same input. On average, sets of inputs

<sup>8</sup>Experimental details and results of the Single-Perspective task on POLISUM are included in Appendix E

Dataset	Domain	Multi-document	Multi-summary	# Gold	Gold Length	% Novel n-grams		
						1	2	3
XSUM	News	✗	✗	11,334	23.3	35.8	83.5	95.5
Multi-News	News	✓	✗	5,622	263.7	17.8	57.1	75.7
CoCoTRIP	Reviews	✓	✓	48 · 3	132.9	22.8	72.4	91.4
POLISUM	Op-Eds	✓	✓	907 · 2	19.7	40.4	86.8	97.5

Table 1: Summary statistics for POLISUM and existing summarization datasets, including data size, reference length, and percentage of novel n-grams in the summary compared to source documents. # Gold for multi-summary datasets is presented as # Samples · # References/Sample. CoCoTrip statistics referenced from Iso et al. (2022).

Headline	Hollywood Strike	
Perspective	Left	Right
Source Passages	...This is a huge issue, but there's another one — AI... A lot of the TV episodes and movies produced by Hollywood are, by nature, highly formulaic...	...By replacing the once majority straight white cast with a ridiculously disproportionate percentage of minorities, industry players alienated a commensurate portion of the audience...
Summaries	The left supports the strikers, arguing that AI will soon be a threat to workers in many industries.	The right argues that Hollywood's troubles stem from its embrace of leftist politics.

Table 2: Example of partial input passages and paired perspective summaries in POLISUM.

contain  $\sim 10$  passages total,  $\sim 5$  passages for each perspective.

Model inputs are constructed by concatenating all passages, ensuring that the left and right perspectives alternate when possible. Due to position biases in summarization (Chhabra et al., 2024), the perspective reflected in the first input document may spuriously impact the resulting summary quality. To mitigate these effects, we report average performance between two variations of the input data. First, models are evaluated on inputs beginning with a left-perspective passage followed by alternating perspectives (e.g., left, right, left, etc.), and then on inputs beginning with a right-perspective passage followed by alternating perspectives (e.g., right, left, right, etc.). The order of individual passages that share a perspective is randomized. Because the reference summaries are limited to a single sentence, all model generations are truncated to the first sentence before evaluation.

### 3.1 Models

A variety of summarization models and instruction-tuned LLMs are considered.

**Extractive Upper Bound.** To represent potential extractive summarization approaches, we estimate an upper bound performance based on semantic similarity to the reference summaries. We consider all sentences contained in the source pas-

sages to be potential extractive summaries. We then select the sentence with the highest BERTSCORE relative to the reference summary and calculate the remaining coverage and faithfulness metrics.

**Neural Summarization Baselines.** To represent neural summarization models, we evaluate the large variants of BART (Lewis et al., 2019), T5 (Raffel et al., 2019), and Flan-T5 (Chung et al., 2022). As BART is not pretrained on any summarization task, we use BART finetuned on the CNN/DailyMail (CNN/DM; Hermann et al. 2015; Nallapati et al. 2016) dataset. In contrast to POLISUM, CNN/DM and other prior summarization datasets lack summaries for individual perspectives. So, in order to encourage BART to generate distinct summaries matching the references' format, the BART decoder is forced to begin generations with "The left" or "The right" respectively. T5 and Flan-T5 are prompted to generate summaries of individual perspectives. See Appendix A for other generation hyperparameter details.

**LLMs.** We evaluate 8 additional LLMs of varying sizes and architectures to establish baselines for current state-of-the-art models<sup>9</sup>. **Mistral-7B** (mistralai/Mistral-7B-Instruct-v0.1; Jiang et al. 2023), **Llama-3.0-8B and 70B** (meta-llama/Meta-Llama-3-8B-Instruct and meta-llama/Meta-Llama-3-70B-Instruct; Dubey et al. 2024), **Llama-3.3-70B** (meta-llama/Meta-Llama-3.3-70B-Instruct), **Vicuna-7B and 13B** (lmsys/vicuna-7b-v1.5 and lmsys/vicuna-13b-v1.5; Chiang et al. 2023), **Mixtral-8x22B** (mistralai/Mixtral-8x22B-Instruct-v0.1; Jiang et al. 2024), and **GPT-4o** (gpt-4o-2024-05-13)<sup>10</sup>. Due to the length of input documents, we evaluate all LLMs in a zero-shot setting. See Appendix B for additional generation hyperparameters, prompts, and reproducibility details.

<sup>9</sup>We use instruction-tuned variants of all applicable LLMs, but for simplicity, omit "Instruct" in names throughout results.

<sup>10</sup><https://openai.com/index/hello-gpt-4o/>

### 3.2 Metrics

**Coverage.** Summary coverage is first measured with ROUGE-1 and ROUGE-L between the generated perspective summaries and their corresponding references (Lin, 2004a). To account for the weaknesses of ROUGE, we also employ BERTSCORE (Zhang et al., 2020) and BLEURT (Sellam et al., 2020) using the *deberta-large-xnli* and *BLEURT-20-D6* checkpoints respectively due to higher correlations with human judgments compared to alternatives<sup>11</sup>. For each metric, we also use a two-sample t-test to determine whether left and right-perspective coverage scores significantly differ ( $p \leq 0.05$ ).

**Faithfulness.** In addition to coverage, we also measure how well the generated summaries are consistent and faithful to the input passages representing the intended perspective. Because faithfulness metrics for our specific task do not exist, we rely on existing factuality metrics as approximations: SUMMAC (Laban et al., 2022) and ALIGNSCORE (Zha et al., 2023). SUMMAC is a NLI-based metric that measures consistency between summaries and source documents. Alternatively, ALIGNSCORE uses information alignment to measure factual consistency. All metrics produce a score between 0 and 100 where higher values reflect more consistent and faithful summaries.

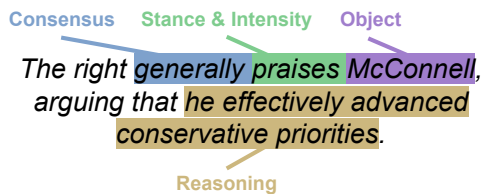


Figure 2: Example perspective summary with indicators of different perspective dimensions highlighted: *Consensus*, *Stance and Intensity*, *Object*, and *Reasoning*

**Perspective.** To measure how well the generated summaries reflect the intended perspective, we decompose dimensions of how perspectives are captured in summaries. *Stance* represents whether the summary accurately portrays the intended perspective’s overall attitude toward the object (e.g., for or against), while *Object* represents whether the summary identifies the salient stance object among input documents. These dimensions correspond to the attitude and stance object of the stance triangle

<sup>11</sup>See human correlations for BERTSCORE linked in [https://github.com/Tiiiger/bert\\_score](https://github.com/Tiiiger/bert_score) and for BLEURT in <https://github.com/google-research/Bleurt/blob/master/checkpoints.md>

theory respectively (Du Bois, 2007). Because our task requires expressing perspectives in natural language as opposed to prior stance detection work, we include two additional dimensions: *Intensity* reflects whether the extent to which a perspective is held or the strength of the belief is accurately captured by a perspective summary (e.g., "praises" rather than "supports" may reflect a stronger belief); and *Reasoning* refers to whether the elaboration in a perspective summary is supported by the input documents. Finally, our task requires summarizing multiple perspectives toward a single object, which may not fully agree. Therefore, we include a final dimension to measure whether perspective summaries appropriately express the extent to which input documents agree or disagree, called *Consensus*.

We conduct a human evaluation of the reference texts and summaries generated by Mistral, Mixtral, and GPT-4o to capture these dimensions. We recruit 6 political science students to rate how well a given summary accurately captures each dimension as reflected in the inputs. In total, each of 40 left and 40 right-leaning summaries randomly sampled from the dataset are judged by 2 annotators.

Additionally, we propose initial metrics to automate selected dimensions drawing on prior work in stance detection and political summarization. To measure *Stance*, we employ JointCL, a zero-shot stance classifier (Liang et al., 2022; Allaway and McKeown, 2023). Using the headlines as the object, we calculate the Root Mean Squared Error (RMSE) between predictions on the generated summaries and references. To capture *Object*, we use a T5-based keyphrase generation model (Li et al., 2023) as a zero-shot stance object classifier, and calculate BERTScore between predictions on generated summaries and input passages. Finally, similar to (Lee et al., 2022), we measure *Intensity* by calculating the average arousal score of polarizing terms according to the VAD lexicon (Mohammad, 2018). Intensity scores are again calculated as the RMSE between arousal scores for generated summaries and input passages. Due to their complexity, we leave further investigation of automated evaluations for *Consensus* and *Reasoning* to future work. See Appendix F and Appendix D for further details on human and automatic metrics respectively.



Model	Left						Right					
	Coverage				Faithfulness		Coverage				Faithfulness	
	R1	RL	BLEURT	BERT	SUMMAC	Align	R1	RL	BLEURT	BERT	SUMMAC	Align
Extractive	13.14	10.25	22.50	55.37	88.85	97.25	13.21	10.34	22.67	55.01	89.23	97.01
BART <sub>CNN/DM</sub>	25.3	22.7	<b>30.5</b>	<b>61.6</b>	24.8	21.0	24.7	22.3	28.9	60.0	25.7	28.9
T5	24.7	22.6	30.9	60.8	23.3	12.0	<b>25.7</b>	<b>23.6</b>	31.2	<b>61.0</b>	24.0	14.7
Flan-T5	24.1	21.2	29.6	58.9	26.2	35.8	24.9	21.8	29.3	58.8	<b>26.9</b>	37.8
Mistral	25.7	21.4	<b>34.5</b>	<b>60.4</b>	25.5	30.1	25.6	21.3	33.2	59.3	26.5	40.5
Mixtral	25.6	21.3	37.2	60.9	24.9	36.9	25.7	21.3	36.3	60.5	25.2	51.7
Vicuna (7B)	24.9	21.6	33.3	59.6	<b>27.9</b>	30.7	25.3	21.8	32.7	58.2	<b>29.4</b>	39.2
Vicuna (13B)	<b>26.4</b>	<b>23.0</b>	37.2	61.4	26.4	37.2	25.3	22.0	35.8	60.3	26.3	<b>54.0</b>
Llama-3.0 (8B)	<b>28.0</b>	<b>23.9</b>	<b>38.1</b>	60.4	25.8	35.2	<b>27.7</b>	<b>23.5</b>	<b>37.3</b>	59.1	<b>26.9</b>	44.8
Llama-3.0 (70B)	23.1	19.0	36.5	59.6	24.0	41.7	<b>23.9</b>	<b>19.6</b>	35.6	59.3	24.4	51.5
Llama-3.3 (70B)	21.6	17.6	35.8	58.8	24.1	<b>42.4</b>	22.2	18.0	34.9	58.2	24.6	<b>51.8</b>
GPT-4o	26.1	21.8	<b>41.8</b>	<b>62.3</b>	24.5	41.0	<b>26.5</b>	21.8	<b>40.9</b>	<b>61.6</b>	24.7	<b>51.8</b>

Table 3: Average coverage and faithfulness scores for the Mixed-Perspective setting. Best model performance on each metric is **bolded** and second best underlined. For each model and metric, significantly higher left than right scores are highlighted in **blue**, and significantly higher right than left scores are highlighted in **red**. Higher scores reflect better performance for all metrics.

Model	Left			Right		
	S↓	O↑	I↓	S↓	O↑	I↓
BART <sub>CNN/DM</sub>	31.0	74.9	15.1	32.6	74.6	14.4
T5	31.1	75.6	17.8	33.7	73.3	18.7
Flan-T5	31.0	77.2	7.9	32.1	73.5	8.0
Mistral	28.8	76.4	8.4	30.7	75.6	6.2
Mixtral	<b>28.2</b>	<b>76.8</b>	10.3	29.4	<b>77.0</b>	5.9
Vicuna (7B)	29.2	76.2	9.0	30.9	75.7	6.6
Vicuna (13B)	29.7	76.7	8.5	31.3	74.9	7.1
Llama-3.0 (8B)	28.9	76.7	8.7	28.9	76.0	6.3
Llama-3.0 (70B)	28.8	76.1	<b>5.8</b>	30.8	75.2	<b>5.3</b>
Llama-3.3 (70B)	29.1	75.8	<b>5.4</b>	30.7	74.3	<b>4.9</b>
GPT-4o	<b>27.5</b>	<b>77.1</b>	6.8	<b>28.3</b>	<b>77.3</b>	<b>5.3</b>
Human Corr.	0.05	0.01	0.09	-0.10	-0.10	-0.02

Table 4: Perspective scores for all models and average sample-level Kendall’s Tau correlation between automatic metrics and human judgments (bottom row). Best performance on each metric is **bolded** and second best underlined. *S*: Stance, *O*: Object, *I*: Intensity.

## 4 Results

### 4.1 Summarization Evaluation

Table 3 presents coverage and faithfulness metrics for all models in the mixed-perspective setting. All models outperform the extractive baseline in coverage, emphasizing that abstractive approaches are vital to performing well on this task. Among abstractive approaches, Llama-3.0 (8B) tends to perform best on ROUGE while GPT-4o tends to perform best on BLEURT and BERTScore. Interestingly, while the larger, 13B-parameter Vicuna model performs better than its 7B-parameter variant, the same is not true for Llama-3.0. This suggests that larger models do not always guarantee better mixed-perspective summarization performance.

Trends in faithfulness metrics differ from that of the coverage metrics. Vicuna models achieve the best performance on SUMMAC for both perspectives and AlignScore for the right-leaning perspective, whereas Llama-3.3 (70B) is best on Align-

Score for the left perspective. Overall, however, all models score far lower on faithfulness metrics than the extractive baseline, suggesting improving model faithfulness to the intended perspective for this task is an important issue to address in future work.

Between perspectives, BLEURT and BERTScore are significantly higher for left-leaning summaries across most models, while SUMMAC and AlignScore are significantly higher for most right-leaning summaries. Notably, all four metrics are model-based as opposed to ROUGE metrics which exhibit few significant differences between left and right-leaning summaries. Because the differences are consistent across most approaches, and prior works have uncovered political biases in models similar to those underlying each metric (e.g., Feng et al. 2023), it suggests that these metrics may also suffer from political biases.

### 4.2 Perspective Evaluation

Model	S	O	I	C	R	Avg. Len	Avg. Ext
Ref.	2.9	3.2	2.9	2.9	1.8	19.7	0.68
Mistral	2.4	3.2	2.5	2.5	1.9	48.6	0.84
Mixtral	3.0	3.4	3.0	3.0	2.6	46.2	0.72
GPT-4o	<b>3.2</b>	<b>3.5</b>	<b>3.2</b>	<b>3.2</b>	<b>2.9</b>	39.7	0.66

Table 5: Average human judgments of each perspective dimension with average length and extraction coverage over all inputs. Best scores are **bolded**. Higher scores reflect better performance for all metrics. Length is measured in Llama-3 tokens. *S*: Stance, *O*: Object, *I*: Intensity, *C*: Consensus, *R*: Reasoning.

Automatic perspective metrics are included in Table 4, and human perspective judgments of Mistral, Mixtral, GPT-4o, and reference summaries are included in Table 5 (breakdown of judgment

2020 Census Battle	
Model	Summary
Reference	The left argues that the new evidence makes clear the GOP’s true motives, and condemns the citizenship question as undermining the purpose of the census.
Mistral	The left argues that the Trump administration’s efforts to include a citizenship question on the 2020 census are a partisan power grab that undermines the accuracy and fairness of the census, and that the real goal is to rig a constitutional mandate.
Mixtral	The left perceives the Trump administration’s attempt to include a citizenship question on the 2020 census as a partisan power grab aimed at benefiting Republicans and non-Hispanic Whites, potentially leading to an undercount of millions of people and an unfair redistricting cycle.
GPT-4o	The left views the Trump administration’s push to include a citizenship question on the 2020 census as a partisan maneuver aimed at disenfranchising minority voters and skewing political power in favor of Republicans and non-Hispanic Whites.

(a) Models often extract different information from individual sources rather than synthesizing the inputs into a single perspective summary.

Amy Coney Barret Confirmed	
Model	Summary
Reference	The right applauds Barret’s confirmation.
Mistral	The right: Republicans did not break any norms in confirming Amy Coney Barret to the Supreme Court.
Mixtral	The right-leaning political perspective within these texts emphasizes that Republicans have not broken any norms or rules in confirming their nominees to the Supreme Court, and that the only rule is having the power to actuate a party’s political will.
GPT-4o	The right argues that Republicans have adhered to historical norms and constitutional principles in their judicial nominations, emphasizing judicial restraint and originalism, while Democrats have engaged in politicizing the process and undermining these norms.

(b) Models tend to extract from the same few input documents.

Table 6: Example perspective summary generations for Mistral, Mixtral, and GPT-4o. Distinct colors signify near paraphrases or exact quotes from a particular document identified manually by the authors.

scores and correlations with automatic metrics are included in Appendix F and Appendix G). GPT-4o summaries tend to score highest on both automatic and human scores, even exceeding reference summary judgments. While the references are similar in extractiveness to GPT-4o, the summaries are significantly shorter which may make them difficult to accurately compare.

Attorney General William Barr Testifies			
Side	Summary	Align	SummaC
Left	The left believes that Attorney General Barr’s handling of the Mueller report was a deliberate attempt to mislead the public and protect President Trump from accountability, and that Congress must take action to hold Trump accountable for his abuses of power.	0.8	22.2
Right	The right believes that Attorney General Barr acted appropriately in releasing a summary of the Mueller report, and that Democrats are unfairly attacking him for not spinning the report in their favor, while also downplaying the fact that Mueller’s report did not find collusion and did not recommend obstruction charges.	94.4	24.4

Table 7: Example pair of left and right-leaning summaries generated by Llama-3.0 (70B) with AlignScore and SUMMAC scores.

Compared to human judgments, the automatic metrics for Stance, Object, and Intensity largely rank the three models in the same order, although there is little sample-level correlation with human judgments. While the metrics may capture system-level differences (i.e., between GPT-4o and Mistral), they often fail to distinguish individual summaries of similar quality (i.e., summaries by the same model). Improving metrics for measuring similarity between perspective dimensions and exploring alternative methods for measuring perspec-

tive dimensions are vital areas for future work.

## 5 What Drives Model Performance?

In this section, we now turn to investigating what factors may be driving performance or present obstacles to high quality summaries. We first qualitatively analyze example summaries, and then analyze the relationship between input features and models’ extraction behavior. In analyzing model extraction behavior, we specifically measure the extent to which models extract from each input document, and what qualities of input documents are associated with higher rates of extraction.

### 5.1 Qualitative Evaluation

Examples summaries generated by Mistral, Mixtral, and GPT-4o are shown in Table 6. In these examples, models appear to attend to the correct perspective within input documents, suggesting that models may have sufficient knowledge of left and right-leaning political beliefs and rhetoric to distinguish inputs in some cases. In summarizing the perspectives, however, models consistently tend to rely primarily on specific information present in a small subset of input documents with little abstractive synthesis of the information (Table 6a). Additionally, models consistently extract or paraphrase from the same documents across many examples, such as in Table 6b. This consistent behavior suggests that qualities of the input documents themselves may drive models’ tendency to extract from them, which we investigate in subsection 5.2.

Furthermore, Table 7 shows a selected pair of left and right summary examples where the right summary faithfulness metrics are notably larger




Figure 3: Average extraction coverage of left and right-leaning source documents for each model shown separately for left and right-leaning summaries. Only the top 4 source documents, 2 for each perspective, are considered. Error bars reflect 95% confidence intervals.

than that of the left. In this case, neither summary contains noticeable faithfulness errors relative to the source documents, but there is a 2 point difference in SummaC scores and a 93 point difference in AlignScore both favoring the right-leaning summary. Throughout the dataset, cases resembling Table 7 where perspective summaries are faithful to similar extents support that neural metrics may suffer from political biases or spurious correlations beyond faithfulness of the input.

## 5.2 Extraction Analyses

To further understand model performance on this task, we analyze extraction behavior with respect to different qualities of the input documents: political perspective, position in the input, length, and use of arousing terms. We use *extractive fragment coverage* to quantify this behavior, as defined by Grusky et al. (2018). In analyses between the sets of left and right-leaning input documents, we hypothesize that models with better performance tend to extract more from the intended political perspective relative to the alternative perspective. Accordingly, we expect less performant models to extract similarly from both perspectives or more from the unintended perspective. For comparisons among documents sharing a common political perspective, we evaluate three additional hypotheses illustrating potentially undesirable behaviors: models extract more information from documents that (1) appear earlier in the input (position bias), (2) are longer (length bias), and (3) that contain more arousing language (arousal bias).

**Document Perspective.** As the Mixed-Perspective setting requires models to disentangle

perspective among input documents, we first analyze whether models tend to extract more tokens from source documents with the correct political slant. To investigate this relationship, Figure 3 presents the extraction coverage scores for each combination of summary and source document perspective (e.g., *left/right* perspective summary and *left/right*-leaning source documents). For both perspectives, performant models tend to extract from source documents portraying the intended perspective (blue bars in the "Left Summaries" plot, and red bars in "Right Summaries") significantly more often than the alternative perspective. The less performant T5, Flan-T5, and Mistral models, however, present insignificant differences in the extent that they extract from the intended perspective with respect to the other perspective. In other words, these models extract similarly from left and right-leaning inputs regardless of the intended perspective of the summary. This suggests that poor performance of the models may be due to the inability to distinguish political leanings among input documents. Uniquely, BART consistently extracts significantly more from right-leaning documents than left-leaning, regardless of the intended perspective. Therefore, **the largest and best performing models (e.g., GPT-4o, Llama-3.3, Llama-3.0, etc.) appear to be sufficiently capable of distinguishing perspectives in inputs.**

**Document Position.** Position and lead biases in summarization systems are well-documented phenomena (Chhabra et al., 2024; Olabisi and Agrawal, 2024; Jung et al., 2019). Following these works,

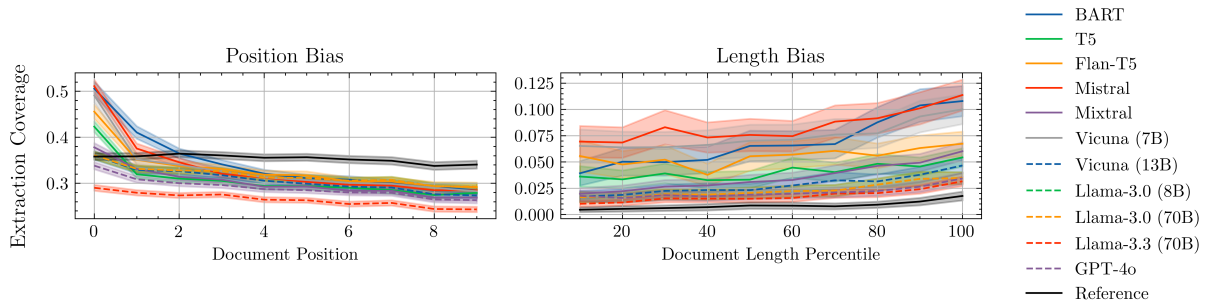


Figure 4: Average extraction coverage of source documents by position in the input (left) and length of input document (right) for each model. For input document length, percentiles buckets are shown rather than raw values. Extraction coverage for length bias analysis only considers terms that appear in a single input document. Smaller numbers reflect earlier input positions. Shaded regions reflect 95% confidence intervals at each document position.




Figure 5: Average extraction coverage of source documents for each model by density of arousing terms in the source document as percentiles.

Figure 4 presents the extraction coverage scores across document positions for each model. As in prior research, we find a near monotonically decreasing relationship between extraction from a document and its position in the input, with all models drawing significantly more terms from the first document compared to those that follow. Among models, Llama-3.3 (70B) and GPT-4o appear to be most robust to document position, indicated by the flatter curves, but still exhibit more bias than human-written references. Surprisingly, however, Mistral appears to be most susceptible to position biases and scores higher on extraction coverage than the smaller BART, T5, and Flan-T5 models. Overall, **all models suffer from position biases and struggle to incorporate documents that appear later in inputs.**

**Document Length.** In addition to position biases, we examine how extraction coverage varies with length of the input documents in Figure 4. As longer documents are more likely to contain terms shared with other documents, in this setting, we modify the extraction coverage metric to only consider terms unique to an individual input document.

In doing so, we more precisely measure to what extent models extract from a particular document in relation to document length. In this case, we see a slight positive trend in extraction of unique terms with document length percentile among models compared to the human-written references<sup>12</sup>. **We find that models tend to extract more unique terms from longer documents, but the impact appears less pronounced than with document position.**

**Document Arousal.** Finally, we consider how the *language* of documents may influence models’ content selection. In particular, we consider the relationship between the density of highly arousing terms in an input document and the extent to which the document is represented in the summary, shown in Figure 5. In contrast to the position and length analyses, the relationship between extraction coverage and arousal of an input document is not consistent across densities of arousing terms. Instead, all models appear to extract less from both non-arousing and highly arousing documents, and to a lesser extent, reference summaries exhibit this behavior as well. This may be because models and humans struggle to incorporate underrepresented perspectives (i.e., documents that present a perspective that diverges further from the other documents in the input). As with other analyses, the largest models appear to be most robust to this trend though still present a similar behavior. Therefore, contrary to our initial hypothesis, we conclude that **models tend to avoid extracting from documents that use extremely many or extremely few arousing terms** relative to other input documents.

<sup>12</sup>For all models excluding Flan-T5, the extraction coverage of the top percentile of document lengths is significantly larger ( $p \leq 0.05$ ) than that of the bottom percentile



## 6 Related Work

**Opinions and Summarization.** Mining and summarizing opinions have long been an interest to the NLP community. Approaches to opinion summarization, typically a multi-document task, have included two-stage extraction-abstraction models (Suhara et al., 2020; Amplayo and Lapata, 2019), pointer-generator networks (Brazinskas et al., 2019; Jayakumar and Malaisamy, 2021), and deep controllable language models (Elsahar et al., 2021). Summarization in other domains, such as news summarization, can also be inherently opinionated. In these cases, prior work has shown that LLMs can fail to fairly represent the diversity of perspectives among inputs (Zhang et al., 2023b; Huang et al., 2023) or account for biases within input texts (Tay, 2019). Furthermore, recent work has found that such behaviors are realized differently when multiple dialects are represented (Olabisi et al., 2022; Olabisi and Agrawal, 2024). In contrast to these works, we focus on a multi-document, multi-summary task with the aim of summarizing political perspectives independently.

**Political Ideologies and LMs.** Recent work has increasingly investigated the political biases and leanings of language models. Santurkar et al. (2023) identify ideologies that language models best align with, while others such as Feng et al. (2023) evaluate the impact of partisan pretraining data on downstream stream task performance. Particularly in summarization, evaluated models have been found to hold name-nationality biases (Ladhak et al., 2023), journalistic framing biases (Lee et al., 2022), geopolitical biases (Li et al., 2024), as well as biases related to perspectives in input documents (Rajan et al., 2023). While our work similarly focuses on a political summarization task, our work uniquely evaluates models’ abilities to summarize individual political perspectives rather than a single bias-free or composite summary.

**Contrastive, Comparative, and Neutral Summarization.** To combat the aforementioned biases in summarization, some have worked toward generating neutral summaries free of journalistic framing bias (Liu et al., 2021; Bang et al., 2023), as well as aligning the stance of the input documents and generated summaries (Liu et al., 2024). Rather than improving the biases of a single summary, a separate line of inquiry has aimed to generate contrasting summaries. Contrastive Opinion Summarization

(COS) was originally dominated by extractive approaches (Kim and Zhai, 2009), selecting a pair of inputs conditioned on each other to be the most contrastive and representative. More recently, work has overcome extractive approach limitations and developed abstractive systems (Iso et al., 2022). These and other abstractive approaches typically rely on knowledge of the perspectives represented by each input. While in our work, we similarly generate independent summary pairs, we instead evaluate how well LLMs’ can summarize *without* prior knowledge of input perspectives.

## 7 Conclusion

In this work, we introduce a dataset, POLISUM, for generating a summary of the left-perspective and a summary of the right-perspective given a set of politically opinionated texts. For this task, we also introduce a framework for evaluating perspective summaries through both human evaluation and automatic metrics. We benchmark current models of varying sizes and architectures on this task, analyzing what behaviors drive performance through their tendency to extract from particular documents.

Models that are capable of accurately and faithfully summarizing multiple perspectives can help alleviate polarization and misunderstandings through exposing others to alternative perspectives than those they hold. While we find that larger and more recent models like GPT-4o can perform well on this task, we also identify problems with faithfulness common among all models evaluated. Additionally, biases due to document position, document length, and use of arousing terms may hinder performance. We recommend future work focus on mitigating these issues while also progressing toward generalization to arbitrary perspectives.

## 8 Limitations

We adopt existing coverage and faithfulness metrics for traditional summarization task as part of our automatic evaluation of summaries. These metrics, however, may not be equally appropriate to the proposed perspective summarization setting. Through additional automatic and human evaluation, we aim to capture aspects of performance specific to summarizing perspectives to complement these results. We encourage future work to investigate the existing coverage and faithfulness metrics on this task, as well as develop novel metrics to ensure accurate measures of performance.

Additionally, while we focus on capturing varied political perspectives, we acknowledge that the left-right dichotomy may not fully capture the broad range of existing political views (Feldman and Johnston, 2014). We restrict our initial evaluation to left and right-leaning perspectives due to data availability, but the task formulation and proposed evaluation framework can be applied to arbitrary sets of perspectives. We encourage future work to evaluate and develop approaches for summarizing a broader range of points-of-view in order to represent such perspective diversity.

Finally, some of the models evaluated may lack appropriate training on recent political issues to enable distinguishing the perspectives presented in source documents. Data in PoliSum covers a time span from September 2018 to January 2024, but the last article of the CNN/DM dataset, for example, was written in April 2015 (Hermann et al., 2015). This lack of recent knowledge may introduce additional challenges to summarizing source passages covering topics such as COVID-19 or recent elections. Future work might explore topic-invariant approaches to summarizing perspectives that are able to adapt to current events.

## 9 Ethics Statement

Hallucinations and factual inconsistencies in language models have received widespread attention due to the risks to safety and potential for misinformation (Huang et al., 2021; Ji et al., 2023). Though the intent of our work is to generate distinct summaries of perspectives in order to better represent each, misrepresentations due to these hallucinations could lead to misunderstandings of opposing perspectives and the accompanying impacts on polarization (Braley et al., 2023; Lees and Cikara, 2021). Additionally, language models that may elevate certain perspectives over others can further contribute to these impacts. We acknowledge that the collected data could be misused in order to intentionally induce these behaviors in language models toward a particular political perspective. In this work, we avoid conducting experiments involving finetuning models and strictly use the collected data to evaluate summaries in order to highlight such weaknesses in existing models. All data in this work is used with permission from The Flip Side and human evaluations included in this work are conducted under an approved IRB protocol.

## Acknowledgments

This work was supported in part by grant IIS-2106666 from the National Science Foundation, the Knight First Amendment Institute at Columbia University, National Science Foundation Graduate Research Fellowship DGE-2036197, the Columbia University Provost Diversity Fellowship, and the Columbia School of Engineering and Applied Sciences Presidential Fellowship. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or the Knight First Amendment Institute. We thank the anonymous reviewers and Smaranda Muresan for providing feedback on an earlier draft of the work.

## References

- Emily Allaway and Kathleen McKeown. 2023. *Zero-shot stance detection: Paradigms and challenges*. *Frontiers in Artificial Intelligence*, 5.
- Reinald Kim Amplayo and Mirella Lapata. 2019. *Informative and controllable opinion summarization*. *arXiv preprint*.
- Stefano Ballelli, Lise Getoor, Daniel G. Goldstein, and Duncan J. Watts. 2021. *Reducing opinion polarization: Effects of exposure to similar people with differing political views*. *Proceedings of the National Academy of Sciences*, 118(52):e2112552118.
- Yejin Bang, Nayeon Lee, and Pascale Fung. 2023. *Mitigating framing bias with polarity minimization loss*. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 11100–11110, Singapore. Association for Computational Linguistics.
- 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, Toronto, Canada. Association for Computational Linguistics.
- Alia Braley, Gabriel S. Lenz, Dhaval Adjudah, Hossein Rahnama, and Alex Pentland. 2023. *Why voters who value democracy participate in democratic backsliding*. *Nature Human Behaviour*, 7(8):1282–1293.
- Arthur Brazinskas, Mirella Lapata, and Ivan Titov. 2019. *Unsupervised multi-document opinion summarization as copycat-review generation*. *CoRR*, abs/1911.02247.
- John R. Chambers, Robert S. Baron, and Mary L. Inman. 2006. *Misperceptions in intergroup conflict: Disagreeing about what we disagree about*. *Psychological Science*, 17(1):38–45.

- Anshuman Chhabra, Hadi Askari, and Prasant Mohapatra. 2024. [Revisiting zero-shot abstractive summarization in the era of large language models from the perspective of position bias](#). *Preprint*, arXiv:2401.01989.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3):6.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#). *Preprint*, arXiv:2210.11416.
- John W Du Bois. 2007. The stance triangle. *Stancetaking in discourse: Subjectivity, evaluation, interaction*, 164(3):139–182.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Hady Elsahar, Maximin Coavoux, Jos Rozen, and Matthias Gallé. 2021. [Self-supervised and controlled multi-document opinion summarization](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1646–1662, Online. Association for Computational Linguistics.
- Alexander R. Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir R. Radev. 2019. [Multi-news: a large-scale multi-document summarization dataset and abstractive hierarchical model](#). *arXiv preprint*.
- Stanley Feldman and Christopher Johnston. 2014. [Understanding the determinants of political ideology: Implications of structural complexity](#). *Political Psychology*, 35(3):337–358.
- Shangbin Feng, Chan Young Park, Yuhan Liu, and Yulia Tsvetkov. 2023. [From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair NLP models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11737–11762, Toronto, Canada. Association for Computational Linguistics.
- Max Grusky, Mor Naaman, and Yoav Artzi. 2018. [Newsroom: A dataset of 1.3 million summaries with diverse extractive strategies](#). 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 708–719, New Orleans, Louisiana. Association for Computational Linguistics.
- Alexander Hanna, Chris Wells, Peter Maurer, Lew Friedland, Dhavan Shah, and Jörg Matthes. 2013. [Partisan alignments and political polarization online: a computational approach to understanding the french and us presidential elections](#). In *Proceedings of the 2nd Workshop on Politics, Elections and Data, PLEAD '13*, page 15–22, New York, NY, USA. Association for Computing Machinery.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. *Advances in neural information processing systems*, 28.
- Nannan Huang, Lin Tian, Haytham Fayek, and Xiuzhen Zhang. 2023. [Examining bias in opinion summarisation through the perspective of opinion diversity](#). In *Proceedings of the 13th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis*, pages 149–161, Toronto, Canada. Association for Computational Linguistics.
- Yichong Huang, Xiachong Feng, Xiaocheng Feng, and Bing Qin. 2021. The factual inconsistency problem in abstractive text summarization: A survey. *arXiv preprint arXiv:2104.14839*.
- Hayate Iso, Xiaolan Wang, Stefanos Angelidis, and Yoshihiko Suhara. 2022. [Comparative opinion summarization via collaborative decoding](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3307–3324, Dublin, Ireland. Association for Computational Linguistics.
- Shobana Jayakumar and Murali Malaisamy. 2021. [Abstractive review summarization based on improved attention mechanism with pointer generator network model](#). *Webology*, 18:77–91.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. [Survey of hallucination in natural language generation](#). *ACM Computing Surveys*, 55(12):1–38.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- 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, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th ophile Gervet, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2024. [Mixtral of experts](#). *Preprint*, arXiv:2401.04088.
- Taehee Jung, Dongyeop Kang, Lucas Mentch, and Edward Hovy. 2019. [Earlier isn't always better: Subaspect analysis on corpus and system biases in summarization](#). *Preprint*, arXiv:1908.11723.
- Hyun Duk Kim and ChengXiang Zhai. 2009. [Generating comparative summaries of contradictory opinions in text](#). In *Proceedings of the 18th ACM Conference on Information and Knowledge Management, CIKM '09*, page 385–394, New York, NY, USA. Association for Computing Machinery.
- 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.
- Faisal Ladhak, Esin Durmus, Mirac Suzgun, Tianyi Zhang, Dan Jurafsky, Kathleen McKeown, and Tatsunori Hashimoto. 2023. [When do pre-training biases propagate to downstream tasks? a case study in text summarization](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3206–3219, Dubrovnik, Croatia. Association for Computational Linguistics.
- 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.
- Jeffrey Lees and Mina Cikara. 2021. [Understanding and combating misperceived polarization](#). *Philosophical Transactions of the Royal Society B: Biological Sciences*, 376(1822):20200143.
- Ro'ee Levy. 2021. Social media, news consumption, and polarization: Evidence from a field experiment. *American economic review*, 111(3):831–870.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. [Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). *arXiv preprint*.
- Bryan Li, Samar Haider, and Chris Callison-Burch. 2024. [This land is Your, My land: Evaluating geopolitical biases in language models](#). *Preprint*, arXiv:2305.14610.
- Yingjie Li, Krishna Garg, and Cornelia Caragea. 2023. [A new direction in stance detection: Target-stance extraction in the wild](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10071–10085, Toronto, Canada. Association for Computational Linguistics.
- Bin Liang, Qinglin Zhu, Xiang Li, Min Yang, Lin Gui, Yulan He, and Ruifeng Xu. 2022. [JointCL: A joint contrastive learning framework for zero-shot stance detection](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 81–91, Dublin, Ireland. Association for Computational Linguistics.
- Chin-Yew Lin. 2004a. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Chin-Yew Lin. 2004b. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Ruibo Liu, Chenyan Jia, Jason Wei, Guangxuan Xu, Lili Wang, and Soroush Vosoughi. 2021. [Mitigating political bias in language models through reinforced calibration](#). *arXiv preprint*.
- Yuhan Liu, Shangbin Feng, Xiaochuang Han, Vidhisha Balachandran, Chan Young Park, Sachin Kumar, and Yulia Tsvetkov. 2024. [P sum: Preserving author's perspective in news summarization with diffusion language models](#). *Preprint*, arXiv:2311.09741.
- Saif Mohammad. 2018. [Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 English words](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 174–184, Melbourne, Australia. Association for Computational Linguistics.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos,  a lar Gu l ehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence RNNs and beyond](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don't give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- Olubusayo Olabisi and Ameeta Agrawal. 2024. [Understanding position bias effects on fairness in social multi-document summarization](#). *Preprint*, arXiv:2405.01790.



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.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2019. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *arXiv preprint*.

Charles Rajan, Nishit Asnani, and Shreya Singh. 2023. [Shaping political discourse using multi-source news summarization](#). *Preprint*, arXiv:2312.11703.

Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cino Lee, Percy Liang, and Tatsunori Hashimoto. 2023. [Whose opinions do language models reflect?](#) *Preprint*, arXiv:2303.17548.

Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.

Yoshihiko Suhara, Xiaolan Wang, Stefanos Angelidis, and Wang-Chiew Tan. 2020. [Opiniondigest: A simple framework for opinion summarization](#). *arXiv preprint*.

Wenyi Tay. 2019. [Not all reviews are equal: Towards addressing reviewer biases for opinion summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 34–42, Florence, Italy. Association for Computational Linguistics.

Thomas Wellings, Regula Hänggli Fricker, Evangelos Pournaras, and Wenqing Fu. 2024. [Brexit on twitter: Unraveling the dynamics of polarization over time](#). In *Proceedings of the 25th Annual International Conference on Digital Government Research*, dg.o ’24, page 855–866, New York, NY, USA. Association for Computing Machinery.

Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. [AlignScore: Evaluating factual consistency with a unified alignment function](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11328–11348, Toronto, Canada. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.

Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B. Hashimoto. 2023a. [Benchmarking large language models for news summarization](#). *Preprint*, arXiv:2301.13848.

Yusen Zhang, Nan Zhang, Yixin Liu, Alexander Fabbri, Junru Liu, Ryo Kamoi, Xiaoxin Lu, Caiming Xiong, Jieyu Zhao, Dragomir Radev, Kathleen McKeown, and Rui Zhang. 2023b. [Fair abstractive summarization of diverse perspectives](#). *Preprint*, arXiv:2311.07884.

## A Zero-Shot Baselines

Models	Prompt
T5, Flan-T5	“summarize the [left,right]-leaning political perspective: {PASSAGES}”
Mistral, Mixtral, Vicuna, Llama-3.0, Llama-3.3, GPT-4o	“Produce a short, single-sentence summary of the [left,right]-leaning political perspective within the following texts. Respond only with the summary beginning with “The [left,right]”: {PASSAGES}”

Table 8: Summarization prompts provided to each model. "left" and "right" are included in the prompt depending on the intended political perspective.

Because it is finetuned on CNN/DM, inference with BART is done without a prompt. BART summaries, however, are forced to begin with either "The left" or "The right" depending on the intended perspective to enforce the task format. Inference with T5 and Flan-T5 is conducted using the first prompt in Table 8. Summaries for all three models are generated with the default transformers package hyperparameters and greedy decoding.

## B Prompting Baselines

For larger LLMs, we use the second prompt in Table 8, generating each perspective summary separately. Larger models are accessed via OpenAI (GPT-4o) and Fireworks AI<sup>13</sup> (Llama-3.3-70B, Llama-3.0-70B, Vicuna-13B, and Mixtral-8x22B) API’s. In all cases, we use a temperature of 0 to ensure deterministic outputs and greedy decoding. Smaller model (Mistral-7B, Llama-3.0-8B, and Vicuna-7B) inference is run with bfloat16 precision, also using greedy decoding.

## C Data Statement

Details characterizing the collected data are included in the following data statement.

### C.1 Curation Rationale

The data was collected in order to evaluate models in summarizing political perspectives on controversial issues provided passages from news articles

<sup>13</sup><https://fireworks.ai>

and op-eds. All data comes from The Flip Side, a site which publishes summaries of left and right-leaning perspectives on current issues with accompanying passages from other news sources. Each sample includes a headline describing the controversial issue, a set of left and right-leaning source passages, and pair of one-sentence left and right perspective summaries.

## C.2 Language Variety

All texts are in English (en-US) and reflect a formal variety of English commonly found in news articles.

## C.3 Speech Situation

All source passages and summaries are originally written texts. As the data is primarily comprised of passages from news articles and op-eds, the texts were intended for a broad American audience, written asynchronously, and are likely edited repeatedly before publishing. All data was originally published between September 2018 and January 2024.

## C.4 Text Characteristics

A majority of the texts in the dataset reflect politically controversial issues, including topics such as elections, policy, political figures, and cultural phenomena. As the task focuses on summarizing political perspectives, the texts also heavily reflect opinions of the original authors in the case of the source passages, or a general left or right-leaning perspective in the case of the reference summaries.

## D Evaluation Metrics Implementation

We evaluate models using existing implementations of coverage, consistency, and factuality metrics. ROUGE is calculated using the ROUGE-score python package (Lin, 2004b). BERTSCORE is calculated using the Torchmetrics implementation<sup>14</sup>. All other metrics are calculated using the implementations by the original authors for each: BLEURT (Sellam et al., 2020), SUMMAC (Laban et al., 2022), and ALIGNSCORE (Zha et al., 2023). All coverage, faithfulness, and automatic metrics for perspective dimensions (Stance, Object, and Intensity) are scaled to a 0-100 range.

<sup>14</sup><https://github.com/Lightning-AI/torchmetrics>

## E Single-Perspective Setting

### E.1 Methods

In the Single-Perspective Setting, models are evaluated similarly to traditional opinion summarization. Models are provided with passages from a single perspective only (i.e., left or right-leaning passages), and asked to generate a summary for the same perspective. Therefore, in contrast to the Mixed-Perspective Setting, models are not expected to distinguish input passages.

### E.2 Results

Automatic coverage and faithfulness metrics for the Single-Perspective Setting are shown in Table 9, while automatic perspective dimension scores are shown in Table 10.

## F Human Judgments

### F.1 Annotator Demographics

Human judgments are conducted with 6 undergraduate students pursuing a BS or BA in Political Science, with self-reported knowledge of the common beliefs of conservatives and liberals. As such, they have the required expertise to judge perspectives conveyed in summaries. Demographics of the recruited annotators are shown in Table 11.

### F.2 Human Judgment Correlations by Model

Correlation with human judgments by model is shown in Table 12. As in the overall sample-level results, correlations with human judgments are largely poor and non-significant. For GPT-4o, however, the automatic metric for Object is significant and weakly negatively correlated with human judgments.

### F.3 Human Judgments by Perspective

Human judgments of perspective dimensions for both the left and right perspectives are shown in Table 13. In both perspectives, GPT-4o summaries are scored highest for all dimensions. While some average scores appear to be slightly higher for the left perspective, no differences between perspectives are significant ( $p > 0.05$ ).

### F.4 Human Evaluation Interface

Screenshots of the human evaluation interface provided to annotators are shown in Figure 6, Figure 7, Figure 8, and Figure 9. Annotators are first provided with instructions describing each perspective

Model	Left						Right					
	Coverage				Faithfulness		Coverage				Faithfulness	
	R1	RL	BLEURT	BERT	SUMMAC	Align	R1	RL	BLEURT	BERT	SUMMAC	Align
BART <sub>CNN/DM</sub>	25.4	22.8	<b>30.7</b>	<b>61.6</b>	25.6	30.3	25.5	22.9	29.9	59.9	27.0	38.4
T5	24.6	22.3	30.1	<b>60.8</b>	23.5	16.3	<b>25.5</b>	<b>23.1</b>	<b>31.2</b>	60.3	23.8	16.8
Flan-T5	23.0	20.2	28.9	59.0	<u>27.3</u>	42.4	23.9	21.2	28.6	58.6	<u>27.3</u>	44.6
Mistral	26.1	22.2	<b>35.9</b>	61.1	26.3	<b>52.1</b>	26.0	22.2	34.7	59.8	27.2	52.7
Mixtral	26.1	22.2	38.6	60.4	26.6	50.7	<b>27.2</b>	<b>23.0</b>	38.0	60.2	26.9	<b>61.4</b>
Vicuna (7B)	24.9	21.5	33.9	<b>59.5</b>	<b>29.6</b>	47.3	26.0	22.4	33.7	58.6	<b>31.7</b>	58.0
Vicuna (13B)	26.0	22.4	<b>38.2</b>	<b>60.3</b>	26.5	51.6	26.3	22.8	37.0	59.4	26.3	<b>60.3</b>
Llama-3.0 (8B)	<b>27.3</b>	<b>23.0</b>	<u>39.2</u>	60.1	25.5	46.0	<b>28.0</b>	<b>23.9</b>	<u>38.9</u>	59.4	25.9	54.3
Llama-3.0 (70B)	23.8	19.6	<u>37.7</u>	60.2	24.2	39.0	25.0	20.5	36.7	59.8	24.4	52.6
Llama-3.3 (70B)	22.2	18.1	<b>37.0</b>	59.1	24.5	42.1	22.8	18.6	35.8	58.5	24.7	53.9
GPT-4o	25.6	21.8	<b>41.3</b>	<b>62.1</b>	25.1	44.2	<u>27.0</u>	<u>23.1</u>	<b>40.8</b>	<b>61.5</b>	25.5	54.2

Table 9: Average coverage and faithfulness scores for the Single-Perspective setting. Best performance on each metric is **bolded** and second best underlined. For each model and metric, significantly higher left than right scores are highlighted in **blue**, and significantly higher right than left scores are highlighted in **red**. Higher scores reflect better performance for all metrics.

Model	Left			Right		
	S↓	O↑	I↓	S↓	O↑	I↓
BART <sub>CNN/DM</sub>	30.3	74.7	15.1	31.7	74.4	14.0
T5	31.3	76.0	17.9	32.5	73.5	19.4
Flan-T5	31.7	<u>77.2</u>	7.9	32.6	73.8	8.0
Mistral	28.5	<b>77.4</b>	8.1	30.0	75.5	6.8
Mixtral	27.4	76.9	9.6	28.9	<b>77.2</b>	5.8
Vicuna (7B)	29.5	76.2	9.1	29.1	75.5	6.5
Vicuna (13B)	<u>27.2</u>	77.0	7.2	31.0	76.3	6.2
Llama-3.0 (8B)	<b>26.7</b>	77.0	7.8	<b>28.5</b>	<u>77.1</u>	6.2
Llama-3.0 (70B)	29.7	77.1	<u>6.1</u>	30.2	77.0	<u>5.1</u>
Llama-3.3 (70B)	27.6	75.3	<b>5.1</b>	30.8	73.3	<b>4.8</b>
GPT-4o	27.6	<b>77.4</b>	6.9	<u>28.6</u>	<u>77.1</u>	5.5

Table 10: Perspective scores for all models in the Single-Perspective setting. Best performance on each metric is **bolded** and second best underlined. S: Stance, O: Object, I: Intensity.

Demographic Variable	Group	%
Gender	Male	66.6%
	Female	33.3%
Political Ideology	Extremely Liberal	16.6%
	Liberal	33.3%
	Slightly Liberal	33.3%
	Moderate	16.6%
	Slightly Conservative	0%
	Conservative	0%
Race/Ethnicity	Extremely Conservative	0%
	American Indian, Alaskan Native, and/or Indigenous	16.6%
	Arab American, Middle Eastern, or North African	0%
	Asian or Asian American	0%
	Black or African American	16.6%
	Latino/a/x or Spanish Origin	16.6%
	Native Hawaiian or Pacific Islander	0%
	Southeast Asian	0%
	White/European American	66.6%
	Other	16.6%

Table 11: Demographic statistics of the 6 annotators.

Model	S	O	I
Mistral	0.04	0.04	0.03
Mixtral	-0.03	0.01	-0.08
GPT-4o	-0.06	-0.19*	0.02

Table 12: Sample-level correlation with human judgments for each perspective dimension by model. Significant correlations are marked with \* ( $p \leq .05$ ).

dimension (Figure 6). Annotators are then shown example summaries with accompanying explanations describing how each perspective dimension could be interpreted in the example summary (Figure 7). We do not include input documents or possible judgments for the example summaries to avoid biasing annotators’ judgments. In the left-hand column, annotators are given the input passages separated by new lines, and in the right-hand column, annotators are provided with a summary followed by questions and 4-point Likert scales for each perspective dimension (Figure 8 & Figure 9). Reference summaries, and summaries generated by Mistral, Mixtral, and GPT-4o are shown on the same page in a randomized order for each set of input passages. Annotators are able to collapse and un-collapse all sections.

## G Coverage and Faithfulness Metric Correlations

Table 14 presents correlations between human judgments on Mistral, Mixtral, and GPT-4o compared to each coverage and faithfulness metric. Among coverage metrics, BLEURT and BERTScore tend to be best correlated among left-perspective summaries, while BLEURT and ROUGE metrics tend to be best correlated among right-perspective summaries. The low correlation coefficients overall

Model	Left							Right						
	S	O	I	C	R	Avg. Len	Avg. Ext	S	O	I	C	R	Avg. Len	Avg. Ext
Ref.	2.9	3.1	2.8	2.9	1.7	19.4	0.65	2.9	3.4	2.9	2.9	2.0	20.0	0.69
Mistral	2.4	3.1	2.5	2.6	2.0	47.1	0.83	2.4	3.2	2.5	2.5	1.8	50.0	0.84
Mixtral	3.2	3.4	3.1	3.2	2.7	45.7	0.71	2.8	3.4	2.8	2.8	2.4	46.8	0.73
GPT-4o	<b>3.3</b>	<b>3.5</b>	<b>3.3</b>	<b>3.3</b>	<b>2.9</b>	40.0	0.65	<b>3.2</b>	<b>3.6</b>	<b>3.2</b>	<b>3.2</b>	<b>2.8</b>	39.5	0.67

Table 13: Average human judgments of each perspective dimension with average length and extraction coverage over all inputs. Scores are shown separately for left and right-perspective summaries. Best scores are **bolded**. Higher scores reflect better performance for all metrics. Length is measured in Llama-3 tokens. *S*: Stance, *O*: Object, *I*: Intensity, *C*: Consensus, *R*: Reasoning.

**Instructions (click to expand/collapse)**

**Task Guidelines**

Thank you for your help with our work on summaries of political perspectives. This task will ask you to read a given text, and rate different aspects of the content and style of the text.

You will be provided with a set of **Passages**, as well as a set of **Summaries**, and will be asked to rate specific attributes of each summary. Passages are separated by newlines and come from different sources.

Rate each summary on the following dimensions:

- Target:** Does the target of the perspective in the summary match that of most input documents?
- Stance:** Does the stance of the perspective in the summary match that of most input documents? Pay special attention to terms such as "supports", "opposes", "agrees", and "criticizes".
- Strength:** Does the intensity or strength of the stance match that of most input documents? Pay special attention to terms such as "skeptical" which may be weaker compared to "criticizes" or "praises" which may be stronger compared to "supports".
- Consensus:** To what extent does the summary indicate agreement or disagreement? Pay special attention to terms such as "divided" and "mixed" as well as qualifiers (or lack thereof) such as "generally".
- Consensus Match:** How well does the summary capture the extent of agreement or disagreement among the *input documents*?
- Reasoning:** Is the reasoning for the perspective provided in the summary supported by most input documents? If no reasoning is given, mark "No Reason".

Example summaries are included below with possible ways to interpret each of the above dimensions, but judgment of how well each are reflected in the input is left to you.

Figure 6: Human evaluation interface screenshot: Instructions.

**Examples (click to expand/collapse)**

**Example 1**

**Summary:** The right generally praises McConnell, arguing that he effectively advanced conservative priorities.

**Explanation:**

In this summary, the target of the perspective is Mitch McConnell, the stance is in support of the target, indicated by praises, the stance is strong indicated by the use of "praises" rather than simply "supports", there is slight disagreement indicated by the use of "generally", and the reason for the right's support is that he effectively advanced conservative priorities.

**Example 2**

**Summary:** The right believes that Biden's efforts to pass voting rights legislation are an overreach of federal power, unconstitutional, and mischaracterize the actual effects of state election laws.

**Explanation:**

In this summary, the target of the perspective is Biden's efforts to pass voting rights legislation, while not explicitly stated, the stance likely opposes the target indicated by the reasoning that it is "unconstitutional", the stance does not appear to be very strong, given by the neutral language in "unconstitutional" and "mischaracterize", there is no use of qualifiers or any indication of disagreement, so the summary indicates consensus agreement, and the reason for the rights criticisms are based in the belief that Biden's actions are unconstitutional among other beliefs.

**Example 3**

**Summary:** The left-leaning political perspective within these texts is that ranked choice voting is a fair and democratic system that allows voters to express their true preferences and prevents the rise of extreme candidates.

**Explanation:**

In this summary, the target of the perspective is ranked choice voting, while not explicitly stated, the stance likely supports the target indicated by the use of terms such as "fair" and "democratic", the stance does not appear to be very strong, indicated by the neutral language in "fair and democratic", there is no use of qualifiers or any indication of disagreement, so the summary indicates consensus agreement, and the reason for the left's supports appears to be that ranked choice voting allows voters to express true preferences and prevents the rise of extreme candidates.

**Example 4**

**Summary:** The right generally praises McConnell, arguing that he effectively advanced conservative priorities.

**Explanation:**

In this summary, the target of the perspective is Mitch McConnell, the stance is in support of the target, indicated by praises, the stance is strong indicated by the use of "praises" rather than simply "supports", there is slight disagreement indicated by the use of "generally", and the reason for the right's support is that he effectively advanced conservative priorities.

Figure 7: Human evaluation interface screenshot: Examples.



### Passages (click to expand/collapse)

**Passages**

"It's important to note that the Texas facility is run not by Border Patrol, but by the Office of Refugee Resettlement [ORR]... In this case, the overflow at ORR is being caused in part by the rise in migrant children arriving at the border alone, not after being separated from parents... "We need to scrutinize whether the administration makes good on its promise to make the conditions under which ORR holds children, including at such warehouse facilities, genuinely more humane. Also crucial is whether the administration undertakes reforms to speed up the process of moving kids from ORR to guardians... [But] Comparing all this to 'kids in cages' confuses the debate in a way that obscures what the Biden administration is genuinely trying to accomplish — and thus makes it harder to actually hold the administration accountable on it."

"Of course, it's never OK to cage children, force them to sleep on floors covered with aluminum-foil blankets and deny them basic sanitation as Trump did. It's never OK to rip babies from their mothers, deport the parent and lose contact with them. Biden isn't doing any of that... "The agency responsible for the temporary care of the unaccompanied children describes the facility as a self-sufficient site equipped with medical and fire services, a dining tent, a soccer field, a basketball court and all sorts of staff, including translators that are available 24/7... Where Biden erred was to open a facility that symbolizes Trump's camps."

"We welcome Biden's more humane approach, but wonder whether it will succeed in the face of the rising tide of juveniles arriving without parents or guardians. Once again the nation is watching its government strain to meet obligations Congress imposed to treat unaccompanied minors with the delicacy they deserve. Once again we see a growing crisis spotlight the broad inadequacies of the government's immigration enforcement system to deal compassionately with human migration... "The solutions require broad vision and actions, including efforts to reduce the instability in Central American countries that send so many people fleeing in the first place. Such efforts, of course, run into a headwind of deep-rooted corruption in some of those countries. But the longer the government leaves those broad solutions unfulfilled, the more it will be forced to deal with waves of migration, one crisis following another. We need a better way of doing this."

"The White House was never likely to convince many Republicans, still in thrall to Mr. Trump, of the merits of a more humane asylum policy, let alone broader goals of legalizing 11 million long-term unauthorized immigrants, or even 'dreamers' brought to the country as children. But a chaotic flood of migrants crossing the border, and the resulting GOP demagoguery about 'catch-and-release,' increases the risk that moderates and independents could recoil from Mr. Biden's push for immigration reform... "The only real solution is a long-term one. It involves sending aid to Central America that improves conditions in the region; beefing up immigration courts with a major infusion of new judges to expand their capacity so that asylum claims are processed quickly and the years-long case backlog is slunk; and, perhaps, screening and processing asylum applicants in Central America, perhaps Guatemala — a prospect that seems far off at best. For the time being, Mr. Biden must manage a balancing act, reestablishing the United States' traditional role as a beacon for beleaguered immigrants while also avoiding a massive new wave of illegal immigration that many Americans would object to."

### Summary 1

**Summary**

The left: Trump's immigration policies were flawed, but the current situation is a crisis that requires a long-term solution involving aid to Central America, increased immigration court capacity, and potentially screening and processing asylum applicants in Central America.

**Perspective Target**

Does the target of the perspective in the summary match that of most input documents?

**1) No Match:** The target is evident in no passages<sup>[1]</sup>
 **2) Slight Match:** The target is covered in a couple passages<sup>[2]</sup>
 **3) Good Match:** The target is covered in most passages<sup>[3]</sup>
 **4) Complete Match:** The target is relevant to all passages<sup>[4]</sup>

**Perspective Stance**

Does the stance of the perspective in the summary match that of most input documents? Pay special attention to terms such as "supports", "opposes", "agrees", and "criticizes".

**1) No Match:** The stance contradicts all passages<sup>[1]</sup>
 **2) Slight Match:** The stance matches a couple passages<sup>[2]</sup>
 **3) Good Match:** The stance matches most passages<sup>[3]</sup>
 **4) Complete Match:** The stance matches all passages<sup>[4]</sup>

**Perspective Strength**

Does the intensity or strength of the stance match that of most input documents? Pay special attention to terms such as "skeptical" which may be weaker compared to "criticizes" or "praises" which may be stronger compared to "supports".

**1) No Match:** The perspective strength matches no passages<sup>[1]</sup>
 **2) Slight Match:** The perspective strength matches a couple passages<sup>[2]</sup>
 **3) Good Match:** The perspective strength matches most passages<sup>[3]</sup>
 **4) Complete Match:** The perspective strength matches all passages<sup>[4]</sup>

Figure 8: Human evaluation interface screenshot: Passages and Questions Part 1.

**Perspective Consensus**

To what extent does the summary indicate agreement or disagreement? Pay special attention to terms such as "divided" and "mixed" as well as qualifiers (or lack thereof) such as "generally".

**1) No Consensus:** Summary indicates no agreement or complete divide<sup>[1]</sup>
 **2) Slight Consensus:** Summary indicates little agreement<sup>[2]</sup>
 **3) Large Consensus:** Summary indicates large agreement<sup>[3]</sup>
 **4) Complete Consensus:** Summary indicates complete agreement or no divide<sup>[4]</sup>

**Perspective Consensus Match**

How well does the summary capture the extent of agreement or disagreement among the passages?

**1) No Match:** The consensus of passages is not represented<sup>[1]</sup>
 **2) Slight Match:** The consensus represents a couple passages<sup>[2]</sup>
 **3) Good Match:** The consensus largely represents passages<sup>[3]</sup>
 **4) Complete Match:** The consensus covers all passages<sup>[4]</sup>

**Perspective Reasoning**

Is the reasoning for the perspective provided in the summary supported by most input documents? If no reasoning is given, mark "No Reason"

**1) No Match:** The reasoning is unsupported<sup>[1]</sup>
 **2) Slight Match:** The reasoning is supported by a couple passages<sup>[2]</sup>
 **3) Good Match:** The reasoning is supported by most passages<sup>[3]</sup>
 **4) Complete Match:** The reasoning is supported by all passages<sup>[4]</sup>
 No Reason<sup>[5]</sup>

Summary 2

Summary 3

Figure 9: Human evaluation interface screenshot: Passages and Questions Part 2.

and particularly for faithfulness metrics, however, suggest that further investigation is needed to adapt coverage and faithfulness metrics to this domain.

## **H Extraction Analyses by Perspective**

In [Figure 10](#), [Figure 11](#), and [Figure 12](#), we present the position, length, and arousal bias analyses by perspective. All identified trends across the left and right perspectives are largely consistent with those reported in the results.

## **I Additional Summary Examples**

Additional summary examples are shown in [Table 15](#), [Table 16](#), and [Table 17](#).

Judge Dimension	Left						Right					
	Coverage				Faithfulness		Coverage				Faithfulness	
	R1	RL	BLEURT	BERT	SUMMAC	Align	R1	RL	BLEURT	BERT	SUMMAC	Align
Stance	-0.03	-0.06	0.22*	0.14*	0.06	0.19*	0.10	0.12	0.24*	0.00	0.02	0.15*
Object	0.02	-0.01	0.25*	0.02	0.08	0.13	0.10	0.13	0.18*	0.11	0.07	-0.01
Intensity	0.05	-0.02	0.27*	0.21*	0.05	0.20*	0.09	0.13	0.26*	0.09	-0.01	0.12
Consensus	0.01	-0.03	0.24*	0.12	0.01	0.18*	0.17*	0.18*	0.27*	0.02	-0.01	0.14*
Reasoning	-0.03	-0.06	0.20*	0.14*	-0.01	0.09	0.08	0.07	0.29*	0.01	-0.00	0.15*

Table 14: Kendall’s Tau correlations between the 5 perspective dimensions measured with human judgments, and the automatic coverage and faithfulness metrics. \* indicates significant correlation ( $p \leq 0.05$ )




Figure 10: Average extraction coverage of source documents for each model by the number of highly arousing terms in the source document (percentiles are shown rather than raw values).




Figure 11: Average extraction coverage of source documents for each model by the document’s position in the input. Lower values reflect earlier positions in the input.

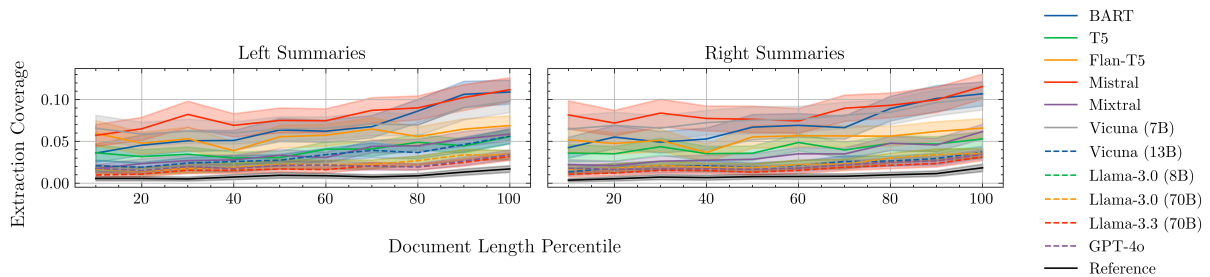


Figure 12: Average extraction coverage of source documents for each model by length of the input document (percentiles are shown rather than raw values).

Ron DeSantis		
Model	Left Summary	Right Summary
Reference	The left argues that DeSantis’s campaign was doomed by miscues and strategic errors.	The right argues that Republican voters were just not ready to move on from Trump.
Mistral	The left: DeSantis’ failure to <i>support aid for Ukraine and his signing of a six-week abortion ban in Florida alienated some moderates and made him look weak, ultimately hurting his chances of winning over voters open to a Trump alternative.</i>	The right: <i>DeSantis’s 2022 re-election campaign was successful</i> , but his 2024 presidential campaign was unsuccessful due to his failure to win over <i>voters who were looking for a fresh start after the Trump-led defeats of 2018, 2020.</i>
Mixtral	The left perceives DeSantis as a weak candidate who failed to effectively differentiate himself from Trump, <i>ultimately losing support from both moderates and evangelicals</i> , and struggling to manage his campaign and image, leading to his downfall in the 2024 GOP presidential primary.	The right believes that DeSantis, despite his initial potential, failed to effectively challenge Trump in the 2024 GOP presidential primary due to his inability <i>to unite a broad coalition, lack of a clear message</i> , and <i>Trump’s political martyrdom following multiple indictments</i>
GPT-4o	The left views DeSantis’ struggles as indicative of the broader Republican Party’s inability to move beyond Trumpism, highlighting internal conflicts, ineffective strategies, and <i>the alienation of moderate voters.</i>	The right views DeSantis’ struggles as a reflection of the Republican base’s enduring loyalty to Trump, despite his flaws, and the challenges of presenting a viable alternative within the party.

Table 15: Additional right-perspective example summaries on the most recent topic in the collected dataset (Ron DeSantis, January 2024). Distinct colors represent quotes or paraphrases from the same document, and italics denote information from passages reflecting the alternate perspective.



Race and Coronavirus		
Model	Left Summary	Right Summary
Reference	The left argues that it is important to focus on racial disparities, because they stem from structural inequality organized along racial lines, and emphasizes the importance of societal solutions.	The right recommends focusing on the factors that underlie racial disparities, such as poverty and pre-existing health conditions, rather than race itself, and emphasizes the importance of individual responsibility.
Mistral	The left: <i>The left argues that the outsized number of deaths in the black community is a serious issue that highlights the nation's failures in serving that community for decades.</i>	The right argues that <i>the outsized number of deaths in the black community is a serious issue that points to the nation's failures in serving that community for decades, but it should not be the focus of the immediate response to the coronavirus.</i>
Mixtral	The left argues that systemic racial and socioeconomic disparities have led to <i>disproportionate exposure to air pollution and health impacts, including COVID-19, among black and Hispanic communities, and that immediate action such as expanding federal financial coverage for testing and treatment, prioritizing hospitals in at-risk neighborhoods, and addressing underlying issues like economic and educational disparities is necessary.</i>	The right perspective suggests that <i>while racial disparities in healthcare are a serious issue, the immediate response to the coronavirus should focus on getting medical resources to where they are needed most, regardless of race.</i>
GPT-4o	The left emphasizes the need to address systemic racial and <i>economic inequalities that exacerbate health disparities</i> , advocating for immediate and <i>long-term measures to ensure equitable access to healthcare and resources for marginalized communities.</i>	The right emphasizes <i>individual responsibility and behavioral factors in health disparities</i> , arguing that immediate pandemic responses should focus on actual risk factors rather than racial generalizations.

Table 16: Additional example summaries for a randomly sampled article. Distinct colors represent quotes or paraphrases from the same document, and italics denote information from passages reflecting the alternate perspective.

Andrew McCabe		
Model	Left Summary	Right Summary
Reference	The left pushes back on the narrative that exercising the 25th amendment would constitute a coup, and criticizes Republicans for hypocrisy regarding the ongoing investigations.	The right condemns McCabe and sees his statements as evidence that the investigation against Trump was politically motivated.
Mistral	The left: <i>McCabe's revelations suggest that Trump's firing of Comey was politically motivated and that Nunes' objections to the FBI investigation may have been driven by bias.</i>	The right: <i>McCabe's revelations suggest that the ongoing FBI investigation of Trump and his campaign may have more to do with politics than the merits of the case, and that some politicians, including Nunes, were aware of the investigation and still worked with the White House to push a conspiracy theory.</i>
Mixtral	The left perceives that top FBI officials, including McCabe, were concerned about <i>Trump's potential threat to national security due to his alleged ties with Russia, and they believe that certain Republican leaders, such as Nunes, McConnell, and Ryan, chose to overlook this concern and defend Trump instead.</i>	The right believes that the FBI, under the leadership of Andrew McCabe and James Comey, <i>overstepped its authority by attempting to act as a check on President Trump, potentially influenced by political bias, and that the use of the 25th Amendment to remove Trump from office was an inappropriate and unrealistic measure.</i>
GPT-4o	The left views <i>McCabe's revelations</i> as evidence that top officials were justifiably concerned about <i>Trump's potential national security risks</i> and that the FBI's actions were a necessary check on presidential power.	The right views the actions of <i>McCabe and other FBI officials as an overreach of their authority, driven by anti-Trump bias</i> , and believes that any attempt to remove Trump from office should follow constitutional processes rather than be decided by unelected officials.

Table 17: Additional example summaries for a randomly sampled article. Distinct colors represent quotes or paraphrases from the same document, and italics denote information from passages reflecting the alternate perspective.