A Narrative Framework for Analyzing Partisan Perspectives in Event Discourse

Anonymous ACL submission

Abstract

Experts from several domains, especially political science, are interested in analyzing political discourse associated with real-world news events. This process would typically require researchers to manually analyze a large collection of news articles on a given event, in order to characterize the underlying partisan perspectives from each side of the political map. Instead, in this work, we propose a systematic approach to summarize partisan perspectives, in an automated manner. Our framework allows us to represent each news article with a predefined structure, comprising of *talking points*, which we then cluster to identify the repeating themes that collectively shape the narrative of an event. Then, we utilize the resulting clusters to generate a summary for each ideology, *left* and *right*, that indicates how each side discusses the event. We show the effectiveness of our framework in capturing partisan perspectives across automated proxy tasks, and human evaluation over a set of events. We release the dataset derived from our narrative framework to the research community.

1 Introduction

One of the signs of the growing social and political polarization is the formation of insulated information bubbles (Gentzkow and Shapiro, 2011; Quattrociocchi et al., 2016; Dubois and Blank, 2018; Garimella et al., 2018), in which news media discourse is shaped around ideological lines, often intended to shape the readers' views on a given event or topic. From a computational standpoint, understanding the differences between ideological perspectives in the coverage of the same news event is typically treated as a straightforward classification problem. Stance and bias detection works (Liu et al., 2022; Luo et al., 2020; Kiesel et al., 2019; Li and Goldwasser, 2019) aim to map news articles to an ideological perspectives, without explicitly explaining the differences. Other works focus on analyzing the discourse, identifying political news framing (Mendelsohn et al., 2021; Field et al., 2018; Card et al., 2015b), sentiment toward relevant entities (Park et al., 2021; Rashkin et al., 2016) or word choices indicative of bias (Recasens et al., 2013).

The view of political discourse understanding as a classification problem has helped provide a straightforward framework for a complex analysis problem, lending itself nicely to the supervised and semi-supervised learning paradigm. Given the rise of Large Language Models (LLM), we suggest that this analysis framework should be revisited. Rather than focusing on specific aspects, such as stance classification, framing or entity sentiment, our goal is to provide a framework that accounts for the interaction between all these elements, the way they contribute to the narrative of a given article, and how their aggregation over multiple articles forms the higher level narratives characterizing the partisan point of view on the news event.

While LLMs have shown a remarkable NL understanding ability, several recent works have demonstrated the challenges LLM face when tasked with analyzing political narratives, either as a straightforward stance prediction task (Ziems et al., 2024), or mapping a given political position to specific stances on various policy issues (Santurkar et al., 2023). We devise our technical approach with these challenges in mind, and suggest a pipeline approach (described in Fig. 1), for extracting the nuanced talking points characterizing the partisan view of a given news event. First, we use an LLM for extracting information from each article using a structured schema, consisting of the main talking points, the relevant entities and the roles they play. Second, we identify repeating talking points by clustering the extracted schema information, and using an LLM to improve the clusters by removing inconsistent and redundant clusters. The goal of this process is to identify repeating topics, which can potentially be shared by both sides of the political map. Finally, we use an LLM to reason about the partisan view of each talking point, capturing the differences in framing and attitudes towards entities expressed by each side.

This process provides several resources, validated through human and automatic validation, that can be used by other researchers. First, we provide a collection of narratives extracted from news articles, organized according to a structured schema. Our dataset consists of $6,141$ articles discussing 24 events related to 4 contested political issues. Second, we identify the prominent talking points characterizing political discourse about a specific event, our data consists of the right- and left-winged perspectives of 389 higher level talking points, relevant for 24 news events. We validate the extracted information in several ways. First, we validate the ability of our framework in characterizing the space of possible talking points through a *topic classification task*. Then, we evaluate the framework's ability to generate partisan perspectives, both automatically using a *partisan classification task*, and through a human evaluation. These results support our final finding, in which we use the extracted representation for stance classification. We show that our approach, extracting abstract partisan talking points, can be used for stance classification over a previously unseen set of documents.

To summarize our main contributions are: **(1)** We propose a new way to conceptualize partisan narrative extraction for news event coverage, which captures nuanced talking points, media frames and entity role analysis. **(2)** We suggest an LLM based pipeline, along with automated validation mechanism, for extracting such partisan narratives. **(3)** we conduct automated and human validation of our pipeline, resulting in a novel dataset, capturing partisan perspective over multiple topics, which can be used to drive future research.

2 Related Work

Prior work on studying partisan perspectives in NLP has primarily focused on *frames*. Framing is a subtle form of media manipulation that highlights some specific aspects of an issue, in order to promote political agendas (Goffman, 1974; Entman, 1993). Card et al. (2015b) proposes the Media Frame Corpus that has 15 generic media frames defined by Boydstun et al. (2014), such as economics or public opinion. In a polarized media environment, frames serve as instrumental mechanisms to promote political agendas through the selective coverage of events (informational bias) and the manipulation of their presentation (lexical bias) (Gentzkow and Shapiro, 2006; Jamieson et al., 2007; Fan et al., 2019). Prior work has also explored approaches to automatically detect and mitigate framing biases. Liu et al. (2019); Akyürek et al. (2020) identify frames through news headlines, Ji and Smith (2017); Khanehzar et al. (2021) detect frames at a document level, and Lee et al. (2022); Liu et al. (2023) mitigate framing biases using multi-domain summarization and graphs.

However, the formalization of frames oversimplifies the intricacies of partisan perspectives and falls short in capturing the nuance of how political agendas are deliberately conveyed in news articles. In this work, look closer at news articles, and represent them with a predefined structure of *talking points*, carefully crafted statements that push the political messages, and cluster them to identify repeating themes, to collectively shape the different partisan perspectives. Recent work has also explored such finer analysis in news articles/political biases. Lawlor and Tolley (2017) presents an entity-focused study of media news framing. Spinde et al. (2021) detects media biases at the word and sentence level, and Frermann et al. (2023) identifies and uses multi-label frames. Our work complements these by introducing a framework that allows us to establish repeating themes of talking points to unveil the partisan perspectives within an event.

3 Narrative Framework

When discussing real world events, political parties and authoritative figures with a relatively large influence typically employ various mechanisms to effectively push their perspectives. In political communication, such carefully crafted messages, statements, or concepts are referred to as *talking points*. These points capture relevant topics with regards to the event, and often push an underlying agenda, that may be partisan biased.

Identifying and understanding these talking points is critical to analyze political discourse surrounding news events. This is because, the prominent talking points for a given event, repeat several times across news articles related to that event, and thus shape the narrative of it. For instance, a right-leaning repeating talking point with respect to the issue `Climate Change` is *Highlighting skepticism towards global climate cooperation, favoring*

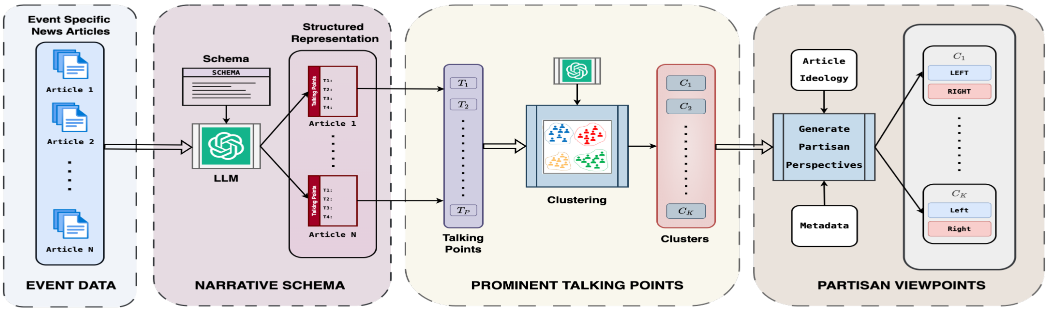

Figure 1: Provides an overview of our narrative framework.

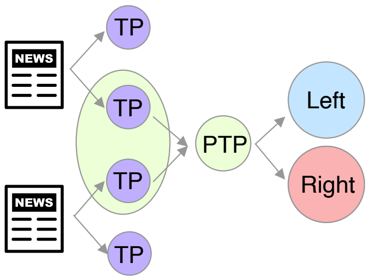

Figure 2: Our Narrative Schema: Given news articles, we extract talking points (TP), which we then cluster (green background), capturing topic relevance. Each cluster is represented by prominent talking point (PTP). Each PTP cluster captures partisan viewpoints for Left/Right political ideologies.

protection of US fossil fuel industries. It clearly captures the right-wing's stance on the issue, and is cardinal in shaping the right narrative.

In this work, we aim to exploit this repeating nature of prominent talking points, in order to summarize the partisan perspectives around an event. To do this, we propose a narrative schema (Fig. 2), that enables us to obtain a structured representation of each news article, that is relevant to the event (Sec. 3.1). We then make use of this structured representation, to analyze the political discourse for the event, and characterize partisan viewpoints indicating how each political ideology discusses the event. Specifically, we use our structured representation to first group the talking points into clusters, which captures their topic similarity. Then, we generate partisan perspectives for each cluster, capturing *left* and *right* political ideologies with respect to the event and a specific topic in it (Sec. 3.2).

3.1 Narrative Schema

We start by identifying the prominent talking points for an event. For this, we propose a schema that analyzes each news article from the event, by defining and building a structure that enables us to summarize the partisan perspectives for that event. Fig. 2 shows an overview. We start with a set of n news articles $\{d_z\}_{z=1}^{n}$, that are relevant to an event $\mathcal{E}$. Our schema reduces every news article d_z to a set of at most four key talking points, i.e. $\{t_i\}_{i=1}^{m}$, where $m \leq 4$. Each point t_i consists of a title and a brief description, explaining the talking point.

Next, we capture additional contextual information around each talking point t_i, by extracting relevant metadata information for it, that could be useful in analyzing partisan differences. To build this metadata, for each point t_i, we first identify the set of entities associated with it. We then capture the relationship between these entities, and how they influence one another, by identifying the set of activities linked to each point t_i. These activities are similar to the form *who did what to whom*. Specifically, each activity consists of a sentence long activity description, an *actor* who is the entity propelling the activity, a *target* that is impacted by the actor, and the sentiment on the target entity, indicating whether the target is positively or negatively impacted by the actor. Finally, we also identify the media frame associated with every activity. The identified media frame follows Boydstan nomenclature (Card et al., 2015a).

In order to obtain the structured representations defined by our schema above, we prompt LLMs. Specifically, we prompt an LLM to identify the key talking points, and the related metadata information, using the prompt template shown in Table 9. Note that we let the LLM decide the key talking points from the news article, although it is possible that it could overlook a prominent talking point. We do this, as we hypothesize that if a talking point is really prominent, then it will repeat in many articles, to shape the narrative. Thus, there is a high chance that the LLM would identify that talking point in other articles, even if the model failed to recognize the prominent point in the given article.

3.2 Characterize Partisan Perspectives

Our overall goal is to analyze political discourse for an event $\mathcal{E}$ by summarizing how each political ideology, say {*left*, *right*}, is talking about the event $\mathcal{E}$. To achieve this, we rely on the schema described in Sec. 3.1, to better characterize the partisan viewpoints for the two political ideologies, rather than directly operating over the news articles.

Specifically, we organize each news article $\{d_z\}_{z=1}^{n} \in \mathcal{E}$ according to the schema, representing every article d_z as a set of at most four talking points, say $\{t_i\}_{i=1}^{m}$, where $m \leq 4$. These talking points are associated only with their respective news article, but we actually want to analyze the partisan discourse for the entire event $\mathcal{E}$. Therefore, we build a talking point set $\mathcal{T} = \{t_s\}_{s=1}^{p}$ for an event $\mathcal{E}$, by aggregating all the four talking points from each article in the event, d_z. Then, to capture the topics of this event, we cluster this set $\mathcal{T}$ to identify topically relevant *prominent talking points*. We utilize the result from the clustering process to generate a left-summary and a right-summary for each cluster, which indicates the partisan viewpoints for the two political ideologies {*left*, *right*}, as it relates to the topics of this cluster. The following describes the prior two steps in more detail.

3.2.1 Clustering the Talking Points

In this first step, we aim to identify a set of *prominent talking points* that are topically relevant to the event, and are sufficient to represent the entire the talking point set $\mathcal{T}$. We do this by grouping topically similar talking points together such that the label associated with each group denotes a *prominent talking point*. Note that the talking points that are clustered together are likely to be topically related, while the cardinality of the cluster indicates the repeating characteristic of the talking point.

Alg. 1 describes the clustering process we use to obtain the *prominent talking points*. First, we embed each point in $\mathcal{T}$ using a dense retriever (Ni et al., 2021) model f, to obtain the corresponding embeddings $\mathcal{Z}$, which we then cluster using the HDBSCAN algorithm, to identify the candidate *clusters*. For each candidate cluster c, we prompt the LLM to characterize the candidate, by generating a cluster label. The cluster label consists of two components, an aspect and a short description about the cluster. The aspect indicates a high-level concept that is discussed in the top-5 talking points, while the description provides a brief summary of the top-5 points. In the prompt, we use the top-5 points closest to the cluster centroid, which we obtain by comparing cosine similarity scores between their respective embeddings.

We note that the output from the clustering process is not entirely perfect, as it is based on traditional distance measures. Therefore, we perform an additional step of updating the cluster label set by merging redundant clusters, and removing inconsistent ones. To remove redundancy, we compare every pair of cluster labels in a greedy manner, and merge the clusters that discuss the same aspect in their cluster labels. The updated label set we obtain after removing redundancy, characterizes the space of possible talking points. More details of the merge and remove process are provided in App. F. The prompt template used to characterize the candidate clusters, and remove redundancy is shown in Tables 22, 23, respectively. Note that the prompts are primarily designed to capture topically relevant talking points.

We then assign each talking point in $\mathcal{T}$ to one of these cluster labels, based on considering the cosine similarity between their corresponding embeddings. This results in a clustering $\{\mathcal{C}_j\}_{j=1}^{k}$ of the talking points along with their associated cluster labels $\{\mathcal{L}_j\}_{j=1}^{k}$, which are termed as the *prominent talking points*. App. F.3 provides further details.

Algorithm 1 Identify prominent talking points

```
Input: Talking points T = {t_s}_{s=1}^p
Initialize: embeddings Z = {z_s = f(t_s)}_{s=1}^p}, n ← no. of news articles, C ← {};
while |Z| > 0.1∗ n do
    clusters ← Clustering(Z);
    labelSet ← [];
    for c in clusters do
        Compute centroid μ_c by averaging;
        Z′ ← getTopKPoints(c, μ_c);
        cLabel ← getClusterLabel(Z′);
        Append cLabel to labelSet;
    end for
    updatedLabels ← updateLabelSet(labelSet);
    S ← TalkingPtMembership(T, updatedLabels)
    T′ ← getClusteredDocs(S)
    T ← T \ T′;
    Z ← Z \ {embeddings of T′};
    Append clusters in S to C
end while
Output: k clusters C = {C_j}_{j=1}^k with cluster labels {L_j}_{j=1}^k
```

3.2.2 Generate Partisan Perspective

Through the clustering process, we have obtained a set of *prominent talking points* that hold topically relevant information. However, these points still do not capture ideology-specific information that is crucial in characterizing partisan perspectives. Therefore, in this step, we provide an ideology label to each talking point in the cluster, and use the metadata associated with the talking point to generate ideology-specific viewpoints for *left*, and *right* political ideologies. These ideology-specific viewpoints indicate how the respective political ideology is discussing the event.

We assign an ideology label {*left, right, center*}, to each talking point in every cluster $\mathcal{C}_j$. We note that every news article d_z gets an ideology label based on its media source. Since each talking point is derived from a news article, it gets the same ideology label as that of the news article.

Next, we describe the process of constructing the partisan perspective of the cluster in detail, by explaining the generation of the left summary (the right summary is generated in a similar manner). Our goal is to generate a summary that it clearly depicts left specific viewpoints, and is different from the right perspective. Therefore, after labeling the talking points in each cluster, we prompt the LLM to generate the left summary in a contrastive manner. For this, we provide the LLM representative talking points from each ideology (left and right), so it can contrast the differences, to identify what defines the left perspective. Specifically, for a cluster $\mathcal{C}_j$, the prompt consists of top-K left-leaning talking points along with top-M right-leaning talking points for contrast. These representative points for each ideology are obtained by considering the cosine similarity between the talking point embeddings and the cluster label embedding.

We observe that context associated with the top-K left-leaning talking points tend to have a myopic localized view of their respective news article. This might not be enough to capture the potential ideological bias exhibited by the article as a whole, which is required to generate a partisan summary. Ideally, the news articles associated with the top-K points should be used to contextualize the prompt. Due to the issue of context length, we resort to working with the news article summaries instead.

We include the news article summaries corresponding to each of the top-K left-leaning points in the prompt. To ensure these article summaries capture the potential ideological bias and topically relevant information in the talking points, we prompt the LLM to generate the article summaries, by conditioning on the ideology label of the article, and the aspect associated with the cluster $\mathcal{C}_j$.

Further, we also want to ensure that the generated left partisan summary captures the relationship between the associated entities in the top-K left-leaning talking points. Therefore, we also include the metadata information consisting of actors, targets, sentiment on the targets, and the relevant media frame as part of the prompt. The prompt template to generate partisan summary is in Table 10.

Issue	No. of Articles	No. of Events
Climate Change	579	8
Capitol Insurrection	1,609	4
Immigration	1,137	4
Coronavirus	2,816	8
Total Count	6,141	24

Table 1: The dataset we use for testing our proposed framework. It is sampled from Nakshatri et al. (2023), and consists of four issues, 6,141 articles, and 26 events.

4 Dataset

To illustrate the effectiveness of our proposed framework, we use the `keyevents` dataset (Nakshatri et al., 2023). This dataset is constructed by segmenting the archive of news articles from NELA-2021 (Horne et al., 2022), into a set of temporally motivated news events. To construct these events, Nakshatri et al. (2023) dynamically analyzed the temporal trend of news articles published for a given issue, and identified the temporal landmarks that could signify the presence of an important news event. Then, the news articles published in and around the temporal landmarks were clustered, to identify all the documents relevant to the news event. In this manner, Nakshatri et al. (2023) proposed a dataset comprising of $40k$ news articles with 611 key news events from 11 issues.

As our goal is to analyze political discourse and characterize partisan perspectives at an event-level granularity, this dataset can be directly applicable to test the efficacy of our narrative framework. Thus, we manually sample a set of six issues and a set of events, which have the highest number of news articles per event, from this dataset. Table 1 shows the detailed statistics of our final dataset.

5 Experiments & Results

Our goal is to evaluate the ability of our framework to generate partisan perspectives. In the following sections, we propose three automated tasks, and a human evaluation. We use ChatGPT [1] as the LLM to analyze each news article, as described in our proposed schema (Sec. 3). Through this, we obtain structured representations for articles from every event in our dataset, and release these, along with the original dataset, to the community.

[1] gpt-3.5-turbo-0125 (OpenAI, 2022)

5.1 Evaluate the prominent talking points

Here, we evaluate our framework's ability to effectively cluster the talking points, using two metrics - *coverage*, and *topic diversity*. A good clustering set should account for most of the documents in the input, while keeping good cluster separation.

To broadly represent all the talking points, in Sec. 3.2.1, we had built a set of *prominent talking points*, by clustering the set of all points $\mathcal{T}$ associated with the event $\mathcal{E}$. These *prominent talking points* capture topically relevant information, for each cluster of points, and as a whole characterize the space of possible talking points for that event.

Coverage. If the *prominent talking points* are actually representative of the cluster, then we expect that each talking point in $\mathcal{T}$ should be able to be mapped back to one of the *prominent talking points*. To evaluate this, we propose a metric called *coverage*, that measures the extent to which the *prominent talking points* collectively capture all points in $\mathcal{T}$ for the event $\mathcal{E}$. Tab. 2 shows the average coverage for each issue. We observe that identified *prominent talking points* cover at least 80% of the talking point set $\mathcal{T}$, indicating they are a good representation of the set $\mathcal{T}$.

Issue	Avg. Coverage per event	Avg. # clusters
Climate Change	83.17	10
Capitol Insurrection	86.70	24
Immigration	90.55	21
Coronavirus	78.18	16

Table 2: Averaged results for coverage.

Topic Diversity. In addition, we also validate if the *prominent talking points* capture diverse topics. Ideally, we want every prominent point to be entirely different from one another, as this would indicate good cluster separation. However, our algorithm merges clusters based on the semantic equivalence of aspects found in the cluster label. Thus, even if the initially formed clusters are far apart in the embedding space, they are merged into a single cluster if they share equivalent aspects in their cluster label. Thus, we cannot use traditional cluster metrics to evaluate diversity.

Task formulation. To this end, we formulate the following ***topic classification task*** : Given a talking point and a set of K' cluster labels, the task is to assign the talking point to the most topically relevant cluster label k^*, where $k^* \in K'$. Note that the talking point is associated with only one of the K' labels, and the rest of the labels are randomly sampled negative examples (other clusters that don't have the talking point). The negative examples help assess the degree of cluster separation. Precisely, k^* helps assess how well the talking point assignments to map to their respective clusters, whereas the remaining negative labels, $K' \setminus \{k^*\}$, help measure the degree of separation between the clusters.

Experimental Setup. To evaluate, we first split the data in each cluster into four quartiles, where the first quartile refers to the top 25% closest talking points to the corresponding cluster label in the embedding space, the 2nd the top 50%, and so on. We randomly sample half the talking points from each quartile for this experiment. We set $|K'| = 4$, i.e. we randomly sample three additional negative labels for each talking point. We prompt ChatGPT to assign the talking point to its most topically relevant label using the prompt in Tab. 11.

Table 3 shows the performance of the topic classification task. We see all quartiles perform well, and documents closer to the cluster label (lower quartile) show strongest topical relevance to the cluster label. This shows that our cluster labels do clearly capture the diverse topics of our talking points, and each cluster captures a unique aspect, when compared to other clusters.

Issue	Q1	Q2	Q3	Q4
Climate Change	91.19	87.47	83.66	80.00
Capitol Insurrection	91.78	89.34	84.56	80.27
Immigration	91.96	88.69	85.01	80.34
Coronavirus	94.07	89.11	84.10	79.94
Avg. Accuracy	**92.74**	**88.90**	**84.37**	**80.12**

Table 3: Averaged results for each quartile for the *topic classification task* indicates that our prominent talking points capture diverse information.

Topic Classification Task + Coverage. Topic classification results indicate that topics associated with *prominent talking points* are diverse ($\approx 80\%$ accuracy) when compared with one another, while the coverage indicate that the prominent points span at least 80% of the set $\mathcal{T}$. On combining both these dimensions, we observe that our approach forms reasonable set of *prominent talking points*.

5.2 Evaluate partisan perspectives

Here, we evaluate the ability of our framework in generating the partisan perspectives.

Partisan. In Sec. 3.2.2, we obtainined ideology-specific viewpoints (summaries) for each cluster $\mathcal{C}_j \in \mathcal{C}$. We now measure the "goodness" of these viewpoints in capturing ideology-specific information. We expect the left-leaning viewpoints will indicate how the *left* political ideology discusses the issue with respect to that cluster, and vice-versa for the right-leaning viewpoints. Therefore, ideally, the left-leaning viewpoints should be entailed by left-biased news articles and should not be entailed by right-biased news articles (similar for right).

Task formulation. To test if the generated ideology-specific viewpoints for each cluster indicate such a partisan behavior, we formulate the following ***partisan classification task***, at the cluster-level granularity. To do it, we use the news article corresponding to each talking point in the cluster.

Given a news article biased towards a particular ideology, say left-biased, and the corresponding left-summary and right-summary for that cluster, the task is to assign the news article to the summary with which it more closely aligns with. In this task, a correct assignment of the news article to its respective summary would indicate that summary exhibits such a partisan behavior.

Experimental Setup. We conduct this experiment in three different settings by changing the manner in which left and right summaries are constructed. In the first setting - *Topically Relevant Points (TPs)*, we construct left summary for a cluster to be the set of 3 topically relevant talking points from left ideology, which are closest to the corresponding left partisan summary for that cluster. Right summary is constructed similarly. In the next setting - *Partisan View*, we construct the summaries using our framework 3.2.2. In the last setting - *Partisan View + Metadata*, in addition to the *partisan view*, we also consider the metadata associated with the entire cluster. Each talking point in the cluster is associated with metadata, which we aggregate for the entire cluster, to facilitate better separation between left and right summaries. More details can be found in appendix C.2. In each case, we prompt the LLM to classify the news article using the prompt template shown in 12. Note that in the prompt, we mask the following terms - *left*-summary, and *right*-summary, and name it summary1 and summary2 instead, so the LLM doesn't classify based on its prior knowledge about *left*/*right* perspectives.

Results. Table 4 shows the overall performance for the partisan classification task (each issue is shown App. A.1). We observe that topically relevant points (TPs) do not clearly distinguish between left and right viewpoints. On the contrary, our partisan view consistently performs better in discriminating between the left and right viewpoints. We also notice that including metadata in the prompt helps improve the performance further.

Approach	Prec.	Recall	F1
TPs	73.44	73.33	73.37
Partisan View	85.03	84.61	84.76
Partisan View + Metadata	**85.93**	**86.14**	**85.98**

Table 4: Averaged results for *partisan classification task* across all issues shows the efficacy of partisan perspectives in capturing ideology-specific information. Tab. 7 shows the breakdown by issue, TPs means Topically Relevant Points.

5.3 Evaluating the usefulness of partisan perspectives

Here, we want to evaluate the generality of the partisan viewpoints. Specifically, we hypothesize that the partisan viewpoints our framework identifies, can be useful to make inferences about *unseen news articles*, that are related to the event.

Task Formulation. To evaluate this, we propose a LLM ***ideology classification task***, at the event-level granularity: Given an *unseen* news article that is related to the event, the task is to predict the ideology of the article by using only the partisan perspectives obtained from our framework.

Experimental Setup. We construct a set of 481 *unseen* news articles that were not part of the initial clustering process, but are closely associated with the events under consideration for each issue. More details on the dataset extraction process is provided in the appendix C.3. As we only know that the article is relevant to the event, to predict the ideology of the news article, we need a partisan summary for the entire event, rather than just each cluster, as our framework builds. Thus, to obtain an event-level partisan summary for an ideology, we concatenate all the summaries from every cluster for that ideology. We then compare the news article embedding with every viewpoint in the concatenated summary set for that ideology, and obtain top-3 closest viewpoints. We call this as *Event Partisan View*.

In addition, for an ideology and for each of the top-3 viewpoints, we also consider the corresponding cluster metadata obtained from 5.2, as it may be able to better distinguish the two political ideologies. We call this as *Event Partisan View + Metadata*. We follow a similar prompting strategy as employed in 5.2, but do not include the ground-truth label for the article ideology as part of the prompt. As a baseline, we directly prompt the LLM to predict the ideology of the given news article. App. 13 shows our prompts.

Results. Table 5 shows the results for this zero-shot LLM ideology classification task. We see that the partisan view outperforms the LLM baseline, and metadata improves performance further. Due to this improvement, we hypothesize that our framework's partisan perspectives are robust enough to capture the potential ideology-specific viewpoints expressed in the news articles, relevant to their events, as they can be used to better detect the ideology of unseen news articles. Further, we also observe that adding metadata can help guide the model to specifically focus on certain entities and their relationships, which are crucial in differentiating the left and right ideological bias. In essence, the results on unseen data indicate that the identified partisan viewpoints are generalizable, and capture repeating themes that shape the narrative for a given political party. More generally, these results show our framework can be used to better analyze a large amount of unseen articles (which has lots of real world uses), as long as they are related to the event initially used in our framework.

Method	Prec.	Recall	F1
ChatGPT	81.38	74.66	73.52
Event Partisan View	77.12	76.64	76.61
Event Partisan View + Metadata	81.20	79.78	79.69

Table 5: Zero-shot ideology classification on unseen news articles, across all issues (see Tab. 8 for each issue). Partisan viewpoints achieve superior zero-shot performance compared to baseline, showing the general benefit of them in improving partisan understanding.

6 Human Evaluation

To better understand and analyze the generated partisan perspectives, we manually annotated data for three randomly sampled events from three different issues (shown in 14). The annotation procedure is described in D, and results below.

6.1 Quantitative Evaluation

We generated partisan perspective for an ideology by leveraging the top-K topically relevant talking points for that ideology, and their respective news article summaries (described in 3.2). Here, we evaluate whether these partisan perspectives for each ideology are actually present in their respective topically relevant talking points, and news articles that were used to construct the partisan viewpoints.

Metrics. We use two metrics. *Summary coherence*, indicates the proportion of partisan summaries that can be constructed from their top-K topically relevant talking points and their associated article summaries. *Mapping quality* (MQ), indicates the proportion of summaries that are actually expressed in the respective news articles.

Table 6 shows the results. We notice a high coherence score for the generated partisan summaries for both political ideologies, implying that the summaries are in agreement with the talking points, and article summaries that were used to construct them. In addition, high mapping quality scores for each ideology indicates that the generated summaries are actually expressed in the news articles.

We manually inspected annotated data, and observed that the generated partisan perspectives are incorrect at times, for example when the LLM fails to produce good news article summaries which are used to generate partisan perspectives. Table 19 shows an example for this. We also notice that the LLM fails at times, to take into account the cited information found in the news articles, which forces the model to generate an incorrect summary.

Issue	L-Coherence	R-Coherence	L-MQ	R-MQ
Climate Change	85.71	100	75.00	76.92
Coronavirus	100	90.90	90.90	70.00
Immigration	93.33	100	84.62	94.44

Table 6: Results from annotated data indicate that partisan perspectives are indeed expressed in the original news articles. **L** refers to left political ideology. **R** refers to right political ideology.

6.1.1 Qualitative results

To further analyze the ideology-specific viewpoints, we ask humans to identify the points of agreement and disagreement between the two partisan summaries. Note that two summaries can disagree only when they discuss broadly the same aspect, and disagree on it. Humans judge this, and Tables 20, 21 shows an example for disagreement and agreement between the partisan perspectives, respectively.

7 Conclusion

In this paper, we proposed a novel approach to capture partisan perspectives for a news event, through a LLM based narrative framework. Our framework captures talking points, both at a topic and partisan level. We conducted automated and human validation on our framework, showing that the partisan perspectives captured by our framework are meaningful and generally useful, such as to improve zero shot ideology classification on unseen news articles. Our future work is to analyze how partisan perspectives dynamically change over time.

8 Limitations and Ethics

To the best of our knowledge, we did not violate any code of ethics when producing this paper. All results are from a Machine Learning model, and should be interpreted as such. We attempted to provide details about our work, both in the main paper and the Appendix, and explain everything thoroughly. In our dataset release, we take care to not release articles that are no longer public.

Our framework itself also has some limitations. As a first, our framework is based on a LLM model, ChatGPT, which is closed source, and the details of unclear. Also, we assume that all the talking points from a left-leaning news source are actually *left-biased*, and vice-versa. However, in reality, it need not be the case. Our approach performs fairly well, even with this assumption primarily because we are only interested in identifying salient talking points from each ideology, and less frequent talking points are rejected.

Our system has many real world applications, but we caution against the safe usage of our framework. Though our approach can be used to identify ideologies, it can also be used in harmful ways, such as users using it to target specific people based on the beliefs that they spread or ideology they align to. For this and meany other reasons, users deploying our work should carefully consider all possible benefits and downsides.

References

Afra Feyza Akyürek, Lei Guo, Randa Elanwar, Prakash Ishwar, Margrit Betke, and Derry Tanti Wijaya. 2020. Multi-label and multilingual news framing analysis. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8614–8624, Online. Association for Computational Linguistics.

Amber E. Boydstun, Dallas Card, Justin H. Gross, Paul Resnick, and Noah A. Smith. 2014. Tracking the development of media frames within and across policy issues.

Dallas Card, Amber Boydstun, Justin H Gross, Philip Resnik, and Noah A Smith. 2015a. The media frames corpus: Annotations of frames across issues. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 438–444.

Dallas Card, Amber E. Boydstun, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015b. The media frames corpus: Annotations of frames across issues. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 438–444, Beijing, China. Association for Computational Linguistics.

Elizabeth Dubois and Grant Blank. 2018. The echo chamber is overstated: the moderating effect of political interest and diverse media. *Information, Communication & Society*, 21(5):729–745.

Robert M. Entman. 1993. Framing: Toward Clarification of a Fractured Paradigm. *Journal of Communication*, 43(4):51–58.

Lisa Fan, Marshall White, Eva Sharma, Ruisi Su, Prafulla Kumar Choubey, Ruihong Huang, and Lu Wang. 2019. In plain sight: Media bias through the lens of factual reporting. 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 6343–6349, Hong Kong, China. Association for Computational Linguistics.

Anjalie Field, Doron Kliger, Shuly Wintner, Jennifer Pan, Dan Jurafsky, and Yulia Tsvetkov. 2018. Framing and agenda-setting in russian news: a computational analysis of intricate political strategies. In *2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Lea Frermann, Jiatong Li, Shima Khanehzar, and Gosia Mikolajczak. 2023. Conflicts, villains, resolutions: Towards models of narrative media framing. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics.

Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2018. Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship. In *Proceedings of the 2018 World Wide Web Conference*, pages 913–922. International World Wide Web Conferences Steering Committee.

Matthew Gentzkow and Jesse M Shapiro. 2006. Media bias and reputation. *Journal of political Economy*, 114(2):280–316.

Matthew Gentzkow and Jesse M Shapiro. 2011. Ideological segregation online and offline. *The Quarterly Journal of Economics*, 126(4):1799–1839.

Erving Goffman. 1974. *Frame analysis: An essay on the organization of experience.* Harvard University Press.

Benjamin Horne, Mauricio Gruppi, and Sibel Adali. 2022. NELA-GT-2021.

Kathleen Hall Jamieson, Bruce W Hardy, and Daniel Romer. 2007. The effectiveness of the press in serving the needs of american democracy.

Yangfeng Ji and Noah A. Smith. 2017. Neural discourse structure for text categorization. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 996–1005, Vancouver, Canada. Association for Computational Linguistics.

Shima Khanehzar, Trevor Cohn, Gosia Mikolajczak, Andrew Turpin, and Lea Frermann. 2021. Framing unpacked: A semi-supervised interpretable multi-view model of media frames. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2154–2166, Online. Association for Computational Linguistics.

Johannes Kiesel, Maria Mestre, Rishabh Shukla, Emmanuel Vincent, Payam Adineh, David Corney, Benno Stein, and Martin Potthast. 2019. SemEval-2019 task 4: Hyperpartisan news detection. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 829–839, Minneapolis, Minnesota, USA. Association for Computational Linguistics.

Andrea Lawlor and Erin Tolley. 2017. Deciding who's legitimate: News media framing of immigrants and refugees. *International Journal of Communication*, 11(0).

Nayeon Lee, Yejin Bang, Tiezheng Yu, Andrea Madotto, and Pascale Fung. 2022. Neus: Neutral multi-news summarization for mitigating framing bias.

Chang Li and Dan Goldwasser. 2019. Encoding social information with graph convolutional networks forPolitical perspective detection in news media. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2594–2604, Florence, Italy. Association for Computational Linguistics.

Siyi Liu, Lei Guo, Kate Mays, Margrit Betke, and Derry Tanti Wijaya. 2019. Detecting frames in news headlines and its application to analyzing news framing trends surrounding U.S. gun violence. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 504–514, Hong Kong, China. Association for Computational Linguistics.

Siyi Liu, Hongming Zhang, Hongwei Wang, Kaiqiang Song, Dan Roth, and Dong Yu. 2023. Open-domain event graph induction for mitigating framing bias.

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*.

Yiwei Luo, Dallas Card, and Dan Jurafsky. 2020. Detecting stance in media on global warming. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3296–3315.

Julia Mendelsohn, Ceren Budak, and David Jurgens. 2021. Modeling framing in immigration discourse on social media. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2219–2263.

Nishanth Nakshatri, Siyi Liu, Sihao Chen, Dan Roth, Dan Goldwasser, and Daniel Hopkins. 2023. Using llm for improving key event discovery: Temporal-guided news stream clustering with event summaries. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4162–4173.

Jianmo Ni, Chen Qu, Jing Lu, Zhuyun Dai, Gustavo Hernández Ábrego, Ji Ma, Vincent Y Zhao, Yi Luan, Keith B Hall, Ming-Wei Chang, et al. 2021. Large dual encoders are generalizable retrievers. *arXiv preprint arXiv:2112.07899*.

OpenAI. 2022. GPT-3.5 (ChatGPT). Computer software.

Kunwoo Park, Zhufeng Pan, and Jungseock Joo. 2021. Who blames or endorses whom? entity-to-entity directed sentiment extraction in news text. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4091–4102.

Walter Quattrociocchi, Antonio Scala, and Cass R Sunstein. 2016. Echo chambers on facebook. *Available at SSRN 2795110*.

Hannah Rashkin, Sameer Singh, and Yejin Choi. 2016. Connotation frames: A data-driven investigation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 311–321.

Marta Recasens, Cristian Danescu-Niculescu-Mizil, and Dan Jurafsky. 2013. Linguistic models for analyzing and detecting biased language. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1650–1659, Sofia, Bulgaria. Association for Computational Linguistics.

Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cinoo Lee, Percy Liang, and Tatsunori Hashimoto. 2023. Whose opinions do language models reflect? In *International Conference on Machine Learning*, pages 29971–30004. PMLR.

Timo Spinde, Lada Rudnitckaia, Jelena Mitrović, Felix Hamborg, Michael Granitzer, Bela Gipp, and Karsten Donnay. 2021. Automated identification of bias inducing words in news articles using linguistic and context-oriented features. *Information Processing Management*, 58(3):102505.

Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. Can large language models transform computational social science? *Computational Linguistics*, pages 1–55.

A Extended Results

A.1 Partisan Classification

Tab. 7 provides the results for partisan classification from Sec. 5.2, across all issues, and the overall performance.

A.2 Zero Shot Ideology Classification

Tab. 8 provides the results for zero shot ideology classification from Sec. 5.3, across all issues, and the overall performance.

B Schema

C Experiments Related

C.1 Prompt Templates for experiment section

C.2 Metadata generation

The goal from this step is to identify the most frequent and discriminative pair of entities along with their sentiments that can help distinguish between the two political ideologies. Note that for an ideology, we use only the top-K documents and its associated metadata for generating the partisan perspective. To account for the metadata from the rest of the members in the cluster and obtain a global cluster-view, we aggregate this information from the top-50% of the members in the cluster. Specifically, we obtain the top-3 target entities that have positive sentiment, and top-3 target entities that have negative sentiment. In each case, we obtain the most common actor associated with the respective target. We also obtain the most common mediaframe associated with the corresponding actor-target context. This information can be plugged into the prompt in addition to the partisan viewpoints to help better distinguish between the political ideologies.

C.3 Dataset extraction

Here, we describe the process used for extracting the set of *unseen* news articles. We note that (Nakshatri et al., 2023) used NELA-2021 dataset for segmenting the news articles into a set of temporally motivated news events. In this process, (Nakshatri et al., 2023) used a temporal window of 3 in order to obtain coherent news events.

In order to obtain *unseen* news articles, yet relevant to the events under consideration, we extend this temporal window to 7 days, and retrieve all the news articles for that time period from NELA-2021 dataset. We filter out all the news articles that part of our clustering process. Then, we consider the all the unseen articles that are closest to the event centroid in the embedding space (threshold ≥ 0.86). Note that we obtain event centroid by averaging the embedding of all news articles relevant to the event. In this way, we extracted 481 relevant news articles for the events under consideration, of which 234 news articles are from right-leaning news sources, and the rest are from the left-leaning news sources.

D Human Evaluation

We conduct human evaluation over a set of three events for three different issues. In this section, we describe the annotation procedure for each task.

Summary Coherence We explain the procedure for *left* political ideology, and the same process is repeated for the *right* ideology as well. First, we explain the task to the annotators with an example. The annotators are provided a left-summary along with three-to-five left talking points and news article article summaries. We ask the annotators to validate if the left-summary can be derived from the news article summaries or the talking points. If it can be derived, then the response is 1, otherwise it is 0. In the cases where annotators are not sure, the response is -1.

Mapping Quality We explain the procedure for *left* political ideology, and the same process is repeated for the *right* ideology as well. In this case, we provide the annotators with a left summary, and a corresponding news article that is most relevant to the left summary (measured based on cosine similarity distance in the embedding space). We segment the news article into sentences of 7, and we only provide the most relevant 7 sentences from the news article to the annotators. First, we let the annotators know that there at at most three points in the left-summary, then ask them to compare the left summary with the left news article content to validate if at least one of the points in the summary is expressed in the article. If it is, then the response is 1, otherwise it is 0. In cases where annotators are not sure, the response is -1.

Qualitative Analysis In this case, we provide a viewpoint from the left-summary and a viewpoint from the right-summary to the annotators, we then ask if the two points are in disagreement with each other. We inform the annotators in prior that the two points must be discussing about the same high-level concept/topic in order for them to be in disagreement with each other. In the cases

Issue	Approach	Avg. Precision	Avg. Recall	Avg. F1-score
Climate Change	Topically Relevant Points	84.11	84.23	84.17
	Partisan View	91.73	89.46	90.29
	Partisan View + Metadata	92.43	90.86	**91.49**
Capitol Insurrection	Topically Relevant Points	69.50	71.62	69.18
	Partisan View	79.33	80.93	**79.93**
	Partisan View + Metadata	81.04	78.08	79.12
Immigration	Topically Relevant Points	69.14	74.64	69.92
	Partisan View	85.38	86.36	85.85
	Partisan View + Metadata	88.27	86.17	**87.15**
Coronavirus	Topically Relevant Points	73.11	72.60	72.77
	Partisan View	83.34	81.76	82.21
	Partisan View + Metadata	83.78	84.20	**83.92**
Overall Performance	Topically Relevant Points	73.44	73.33	73.37
	Partisan View	85.03	84.61	84.76
	Partisan View + Metadata	85.93	86.14	**85.98**

Table 7: Averaged results for *partisan classification task* shows the efficacy of partisan perspectives in capturing ideology-specific information.

Issue	Method	Avg. Precision	Avg. Recall	Avg. F1-score
Climate Change	Zero-shot chatGPT	82.47	71.29	70.83
	Event Partisan View	76.21	75.36	75.60
	Event Partisan View + Metadata	80.55	76.55	**77.01**
Capitol Insurrection	Zero-shot chatGPT	80.24	72.14	74.19
	Event Partisan View	80.00	83.75	80.91
	Event Partisan View + Metadata	83.82	84.91	**84.32**
Immigration	Zero-shot chatGPT	76.81	76.54	70.94
	Event Partisan View	80.14	82.13	78.47
	Event Partisan View + Metadata	81.73	83.56	**82.21**
Coronavirus	Zero-shot chatGPT	84.81	82.05	**81.19**
	Event Partisan View	69.16	68.96	68.96
	Event Partisan View + Metadata	77.29	76.83	76.48

Table 8: Compares the performance of zero-shot ideology classification on unseen news articles. On average, partisan viewpoints are able to achieve good zero-shot performance compared to baseline.

Schema	Schema Prompt Template (incremental)
Talking Point	You are tasked with discerning the key talking points from the following 'NEWS ARTICLE'. Your objective is to condense the contents of the 'NEWS ARTICLE' into a succinct list of up to four primary talking points, each accompanied by a brief description. \n 'NEWS ARTICLE': {article}
Entities	For each 'Talking Point', identify up to three pivotal entities associated with it, and assess whether the author of the 'NEWS ARTICLE' exhibits a bias either against or in favor of the mentioned entities. In instances where no discernible bias is evident, categorize the entity as neutral. The goal is to compile a list of entities along with their entity types, categorizing them as either against, in favor, or neutral, and accompany each categorization with a brief explanation.
Activity	For each 'Talking Point' and its associated 'Entities', identify the primary activities linked to it. For every identified 'Activity', pinpoint the entity assuming the role of 'Actor', driving the said 'Activity', and the entity acting as the 'Target', which is influenced by the 'Actor'. Assess whether the impact on the 'Target' is positive, negative, or neutral, providing a rationale for the impact. Focus only on pivotal 'Activities' closely related to the 'Talking Point'.
Media Frame	For each 'Talking Point' and its corresponding 'Activity', predict its media frame, and categorize it into one of 15 labels: Economic, Capacity & Resources, Morality, Fairness & Equality, Legality, Policy, Crime, Security, Health, Quality of Life, Cultural, Public Opinion, Political, External Regulation, or Other. With respect to the predicted 'Frame', provide a short explanation on how it is related to the main 'Activity'.

Table 9: Prompt template used to obtain the structured representation of the article along with the relevant metadata.

when the annotator finds that there is a disagreement, the annotators are asked label 1; otherwise 0. In the cases where the annotators are not sure, the response is -1. The same setup is replicated to validate if there is an agreement between the two points.

Note that our annotators were computer science Ph.D. students who were under the age of 30.

Generate Partisan Summary

You are provided with an aspect of discussion related to a news event, along with biased talking points from left and right political ideologies discussing the same aspect. Each talking point is associated with its respective news article summary, and metadata that includes actions, actors, targets, impacts, and framing. On comparing and analyzing the talking points from both ideologies, the objective is to refine and condense left-biased talking points into at most three unique points, such that the new points clearly capture the political bias towards left ideology. Redundant points and those not aligning well with left political ideology should be excluded.
Input
Aspect of discussion: {aspect}
Left-biased talking points: {left-biased points} ## includes metadata for each point
Left-biased news article summaries: {left-biased summaries}
Right-biased talking points: {right-biased points}

Table 10: Prompt template used to obtain the partisan summary for left political ideology. Similar prompt is used for the obtaining partisan summary for the right political ideology as well.

Topic Classification Task

You are given a 'document' and four labels, all derived from the same news event. The task is to determine the most topically relevant label to the document. Your goal is to assign the document to only one of the four labels. If the document is topically relevant to 'label1', please respond 'label1'. If the document is topically relevant to 'label2', please respond 'label2'. If the document is topically relevant to 'label3', please respond 'label3'. If the document is topically relevant to 'label4', please respond 'label4'. Strictly refrain from providing additional information.
Input to analyze
'Document': {doc}
'Label1': {lab1}
'Label2': {lab2}
'Label3': {lab3}
'Label4': {lab4}

Table 11: Prompt template used for the topic classification task.

Partisan Classification Task

Given a segment of a 'news article' from a {ideology}-biased media source and two summaries derived from the same news event, your task is to perform binary classification by assigning the news article to one of the two summaries. Each summary has a set of talking points about the event. Each summary is also accompanied by metadata that includes frequently occurring actors, targets, sentiment on the target entities, and media frame associated with the context of the talking points. Your goal is to use the associated metadata to better determine if the provided news article segment has a viewpoint that is more similar to 'summary1' or 'summary2'. The response should strictly be 'summary1' when the 'news article' segment has a consistent viewpoint with 'summary1'; otherwise, it should be 'summary2' indicating the 'news article' has consistent viewpoint with 'summary2'.
Refrain from providing any additional information.
Input to analyze
'News article': {article}
'Summary1':{summ1}
'Summary1 Metadata':{summMetadata1}
'Summary2':{summ2}
'Summary2 Metadata':{summMetadata2}

Table 12: Prompt template used for the partisan classification task.

Ideology Classification Task (Baseline Prompt)

The task is to perform a binary classification to determine whether the ideology of the given 'news article' leans more towards the 'left' or the 'right'. You are to output one of the two labels. Strictly adhere to the following output format, and refrain from providing additional information.

'News Article': {articleContent}

Table 13: Prompt template used for the zero-shot ideology classification task (baseline).

E Examples

F Clustering the talking points

As described in 3.2.1, we cluster the initial talking point set to identify the prominent talking points. In this process, we merge redundant clusters and remove incoherent clusters. The details of this process is outlined in this section.

F.1 Merge Redundant Clusters

In order to merge redundant clusters, we devise a simple greedy algorithm. We consider pairwise combinations of cluster labels, and prompt the LLM to verify if the clusters can be merged based on the prompt template shown in 23.

We start by constructing a set of pairwise cluster labels $\mathcal{S} = \{(s_1, s_2), \cdots, (s_{n-1}, s_n)\}$. Precisely, for every cluster, we consider all possible pairwise combinations for the top-7 closest labels to that

Issue	News Event
Climate Change	**Event Title:** Biden Announces Ambitious Greenhouse Gas Emissions Cut **Event Description:** This is about President Joe Biden's announcement of an ambitious cut in greenhouse gas emissions as he looks to put the US back at the center of the global effort to address the climate crisis and curb carbon emissions.
Coronavirus	**Event Title:** Biden's COVID-19 Vaccination Mandate **Event Description:** This is about President Joe Biden's announcement of new COVID-19 vaccination requirements for federal government employees, healthcare workers, and companies with 100 or more employees, and his criticism of politicians who are undermining trust in COVID vaccines.
Immigration	**Event Title:** Biden's Refugee Cap Decision **Event Description:** This is about the criticism faced by President Biden for his decision to not raise the cap on refugees allowed to enter the US this year, which he had promised to do during his campaign.

Table 14: Events considered for human evaluation.

	Structured Representation from article
News Article	(CNN) A White House riding high on a wave of ambition is setting up a series of inevitable tests of whether Joe Biden is promising more than he can deliver . The President 's aggressive pledge to cut US carbon emissions unveiled at his global online summit Thursday is the latest audacious bet in a presidency that is notable for a moderate tone but an increasingly expansive progressive agenda .\nThe scale of Biden 's plans that he will try to sell to the nation in an address to Congress marking his first 100 days next week shows @ @ @ @ @ @ @ his power to forge a legacy as a generational reformer .\nSome admirers have already put him in the company of great Democratic Presidents like Franklin Roosevelt and Lyndon Johnson .\nBut for such praise to be accurate , he will have to pass an agenda that aims to overhaul much of the economy to benefit American workers – and he 's about to hit a wall of Republican opposition in divided Washington .\nSo while the President has enjoyed a fast start , remarkable success in accelerating vaccines to fight the pandemic @ @ @ @ @ @ @ stirring , real questions are mounting over his capacity to follow through .\nA treacherous road lies ahead that will require Biden to convince the public to embrace all of his programs and to make his opponents pay a price for opposing them .\nThat 's one reason why Biden 's remarks opening a climate summit that included leaders like China 's Xi Jinping and Russia 's Vladimir Putin sounded more like a speech in a Pittsburgh union hall than the blueprint of a leader bent on a costly crusade to save the @ @ @ @ @ @ @ climate , I think jobs , " Biden said , billing the fight against global warming as an extraordinary economic opportunity that will put Americans to work capping abandoned oil wells and assembling electric cars . ...
Talking Point 1	Title: Biden's ambitious climate pledge Description: Biden unveiled an aggressive plan to cut US carbon emissions at a global online summit. The scale of his plans shows his power to forge a legacy as a generational reformer, but questions are mounting over his capacity to follow through.
Talking Point 2	Title: Republican opposition and challenges ahead Description: Biden is likely to face significant Republican opposition in his efforts to pass his agenda, including his plans to overhaul the economy and address climate change. He will need to convince the public to embrace his programs and make his opponents pay a price for opposing them.
Talking Point 3	Title: Biden's broader policy agenda Description: In addition to climate change, Biden has voiced support for sweeping election reform and infrastructure plans. Passing these bills is seen as necessary to preserve the party's hopes in future federal elections and to address American ideals about equal access to the franchise.
Talking Point 4	Title: The difficulty of compromise and the need for Democratic unity Description: Biden's aspirations may face challenges due to the lack of compromise in modern politics and the potential for Republican obstruction. Biden's unwillingness to pare down his aspirations and accept compromises may make it difficult to achieve his goals.

Table 15: Shows the reduction of the news article to its respective talking points.

	Prominent Talking Point Generation
Prominent Point	Key aspect: Opposition and Challenges to Biden's Climate Change Agenda Summary Description: The articles collectively highlight the significant opposition and challenges President Biden faces in pushing forward his ambitious climate change agenda. Republican resistance, concerns about economic impact on traditional energy sectors, difficulties in securing funding and political support, and obstacles in translating rhetoric into action are key themes discussed in relation to Biden's climate initiatives.
Talking Point 1	Title: Obstacles and opposition to Biden's climate change agenda Description: Republicans have vowed to fight against Biden's proposals to shift the U.S. energy sector away from fossil fuels, indicating potential challenges at home.
Talking Point 2	Title: Political challenges and opposition Description: The article mentions the challenges Biden faces in keeping political support and securing funding for his ambitious climate goals. It also highlights Republican opposition, arguing that transitioning to clean energy would harm American oil, natural gas, and coal workers.
Talking Point 3	Title: Republican opposition and challenges ahead Description: Biden is likely to face significant Republican opposition in his efforts to pass his agenda, including his plans to overhaul the economy and address climate change. He will need to convince the public to embrace his programs and make his opponents pay a price for opposing them.
Talking Point 4	Title: Challenges and opposition Description: The article highlights the difficulties Biden and his team may face in converting their bold rhetoric into action. It mentions potential obstacles such as the fate of Biden's infrastructure plan, Republican opposition to climate initiatives, and the power of the Supreme Court to strike down laws limiting carbon pollution.
Talking Point 5	Title: Challenges in passing Biden's agenda Description: The article mentions that Biden will face opposition from Republicans in Washington, which could pose a challenge to passing his agenda. Questions are raised about Biden's capacity to follow through on his plans, particularly in overhauling the economy to benefit American workers.

Table 16: An example showing a topically relevant prominent talking point that is constructed using top-5 talking points shown.

cluster in the embedding space. For each element in $\mathcal{S}$, we prompt LLM to infer if the pair of labels are discussing about the same aspect. If the aspects, say (s_1, s_2), are equivalent, then we merge these aspects, and update the set $\mathcal{S}$ by removing every element in the set that contains s_1 or s_2. In the second iteration, we construct a new set, $\mathcal{S}'$, that holds every combination of updated cluster labels,

Partisan Viewpoints	
Prominent Point	Key aspect: Opposition and Challenges to Biden's Climate Change Agenda Summary Description: The articles collectively highlight the significant opposition and challenges President Biden faces in pushing forward his ambitious climate change agenda. Republican resistance, concerns about economic impact on traditional energy sectors, difficulties in securing funding and political support, and obstacles in translating rhetoric into action are key themes discussed in relation to Biden's climate initiatives.
Left Talking Points	1. Title: Republican opposition and challenges ahead Description: Biden is likely to face significant Republican opposition in his efforts to pass his agenda, including his plans to overhaul the economy and address climate change. He will need to convince the public to embrace his programs and make his opponents pay a price for opposing them. 2. Title: Challenges in passing Biden's agenda Description: The article mentions that Biden will face opposition from Republicans in Washington, which could pose a challenge to passing his agenda. Questions are raised about Biden's capacity to follow through on his plans, particularly in overhauling the economy to benefit American workers. 3. Title: Climate change has become a centerpiece of President Biden's economic agenda Description: Over the past few years, addressing climate change has shifted from a backburner issue to a crucial part of President Biden's domestic agenda and economic policy. 4. Title: Republican opposition and challenges ahead Description: Biden is about to face a wall of Republican opposition in divided Washington. He will need to convince the public to embrace his programs and make his opponents pay a price for opposing them. 5. Title: Climate as a centerpiece of Biden's economic agenda Description: Biden has made climate a central focus of his economic agenda, seeing it as an opportunity for job creation and economic growth. His plans include transitioning to clean energy, promoting electric vehicles, and investing in infrastructure.
Right Talking Points	1. Title: Economic Concerns and Job Losses in Fossil Fuel Industries Description: Right-leaning critics express worries about job losses and economic impacts on industries like oil, natural gas, and coal due to Biden's climate agenda, contrasting it with the potential benefits of transitioning to cleaner energy sources. 2. Title: Opposition to Biden's Climate Policies Description: Republicans and conservative voices, including the Wall Street Journal Editorial Board, criticize Biden's climate policies as detrimental, labeling them as a disguised 'Green New Deal' that could lead to job destruction in fossil fuel and carbon-intensive sectors. 3. Title: Concerns Over Funding and Tax Increases Description: Right-leaning perspectives highlight concerns about the funding sources for Biden's climate initiatives, particularly criticizing potential tax increases on corporations and the economic challenges associated with financing the climate agenda.
Left Perspective	L1: Title: Republican Opposition to Biden's Climate Agenda Description: Biden's climate change agenda faces significant challenges from Republican opposition, highlighting the partisan divide on environmental policies and the struggle for bipartisan support in a politically divided landscape. L2: Title: Biden's Climate Agenda as Economic Focus Description: Climate change has become a central component of President Biden's economic agenda, emphasizing the integration of environmental goals with job creation and economic growth, reflecting a progressive approach to tackling climate issues. L3: Title: Challenges in Passing Progressive Climate Legislation Description: Biden's plans to overhaul the economy and address climate change are met with skepticism regarding their feasibility and potential impact on American workers, underscoring the need for navigating opposition and ensuring successful implementation of progressive climate policies.
Right Perspective	R1: Title: Economic Concerns and Job Losses in Fossil Fuel Industries Description: Right-leaning critics express worries about job losses and economic impacts on industries like oil, natural gas, and coal due to Biden's climate agenda, contrasting it with the potential benefits of transitioning to cleaner energy sources. R2: Title: Opposition to Biden's Climate Policies Description: Republicans and conservative voices, including the Wall Street Journal Editorial Board, criticize Biden's climate policies as detrimental, labeling them as a disguised 'Green New Deal' that could lead to job destruction in fossil fuel and carbon-intensive sectors. R3: Title: Concerns Over Funding and Tax Increases Description: Right-leaning perspectives highlight concerns about the funding sources for Biden's climate initiatives, particularly criticizing potential tax increases on corporations and the economic challenges associated with financing the climate agenda.

Table 17: Shows an example of partisan summaries obtained for a cluster.

and repeat the previous step. We run the algorithm for two iterations or halt if there are no merges after the first iteration. Considering the cost constraints associated with chatGPT, we consider top-7 closest cluster labels,

Coherent Example
Right Partisan Summary: R1: Title: Opposition to Biden's Climate Goals Description: Right-leaning sources criticize Biden's emission reduction targets, highlighting concerns over economic costs, job losses, and potential negative impacts on industries like American oil and automobile sectors. R2: Title: Skepticism Towards Clean Energy Investment Description: Republicans express skepticism towards Biden's plans for massive investment in clean energy technologies, raising concerns about the associated costs, tax increases, and economic impact on American workers. R3: Title: Critique of Lack of Implementation Details Description: Right-leaning articles criticize the lack of specific details provided about how emission cuts will be achieved, highlighting concerns about economic damage, job losses, and the transparency of the implementation plans.
Topically relevant right talking points: 1. Title: Far-reaching changes required to meet emission reduction goals Description: Achieving a 50% reduction in emissions by 2030 would require significant changes, such as increasing renewable energy sources, transitioning to electric vehicles, shutting down coal plants, and adopting new energy efficiency targets in industries. 2. Title: Investment in clean energy Description: The summit highlighted the case for massive investment in clean energy technologies and infrastructure, both in the US and around the world. This investment is seen as crucial for creating prosperous and cleaner economies in the long run. 3. Title: Funding for carbon capture technology and critical minerals Description: The Department of Energy will begin accepting applications for a 75 million fund for carbon capture and storage technology. Additionally, 19.5 million in awards will be available for extracting critical minerals used in developing batteries and components for electric vehicles. 4. Title: Lack of details and economic cost Description: The article criticizes the lack of details provided about how the emissions cuts will be achieved and the potential cost to industries and American consumers. It suggests that the economic damage caused by the plans is conveniently ignored by the media.
Corresponding news article summaries: 1. President Biden's ambitious pledge to cut emissions by 2030, including substantial financial support for developing countries, is portrayed as a challenging and costly endeavor by a right-leaning source. The article emphasizes the significant economic changes required to achieve these emission reduction goals, highlighting the potential need for government subsidies and carbon taxes. It also underscores the limited global impact of U.S. emissions reductions compared to major emitters like China and India, suggesting that the efforts may not be effective in combating climate change on a global scale. 2. Biden closes global climate summit emphasizing collective action and massive investments in clean energy, contrasting with Republican skepticism and focus on economic impact on American workers. The article highlights Biden's push for clean-energy technology and infrastructure funding, facing opposition from Republicans who criticize the cost and potential tax increases, while emphasizing the economic benefits and job creation potential of transitioning to clean energy. 3. President Biden's administration targets job creation in fossil fuel communities through federal investments in renewable energy, aiming to reduce coal and petroleum emissions. The report outlines funding for carbon capture technology, critical minerals extraction, rural broadband expansion, and infrastructure projects. While some Republican leaders in coal-dependent states like West Virginia oppose Biden's climate agenda, others, including the nation's largest coal miners' union, support transitioning away from fossil fuels for a cleaner energy future. 4. Right-leaning news article criticizes left-wing media's praise of President Biden's climate plan as 'ambitious' and 'visionary', highlighting concerns about economic damage and lack of details on implementation and costs, while emphasizing the partisan bias in media coverage of climate policies.

Table 18: Shows a coherent example. The generated partisan summary is correct.

Incoherent Example
Right Partisan Summary: R1: Title: Criticism of Biden's vaccine mandate as dictatorial Description: Dr. Fauci's characterization of President Biden's vaccine mandate as 'moderate' is challenged by right-leaning perspectives, labeling the mandate as dictatorial and overly lenient, emphasizing concerns about individual liberty being compromised. R2: Title: Opposition to perceived leniency in vaccine mandate Description: Right-leaning voices critique Dr. Fauci and President Biden for what they view as a lenient approach in the vaccine mandate, suggesting a need for stricter measures to ensure compliance and public health protection. R3: Title: Advocating for a stricter vaccination-only policy Description: Right-leaning perspectives highlight a preference for a stricter vaccination-only policy over the compromise of allowing testing as an alternative in President Biden's vaccine mandate, emphasizing the importance of prioritizing vaccination for public health benefits.
Topically relevant right talking points: 1. Title: President Biden's vaccine mandate is considered somewhat 'moderate' by Dr. Anthony Fauci Description: Dr. Fauci believes that President Biden's vaccine mandate strikes a balance between the need for vaccination and the concerns of those who do not want to get vaccinated but also do not want to lose their jobs. 2. Title: Dr. Fauci's stance on the vaccine mandate Description: Dr. Fauci is quoted as saying that the President's approach is somewhat moderate and a compromise, but the article disagrees, stating that Fauci's true stance on vaccine mandates is even less moderate. 3. Title: President Biden's vaccine mandates Description: Dr. Fauci describes President Biden's vaccine mandates as a moderate approach, allowing for testing as an alternative to vaccination for those who are unwilling to get vaccinated.
Corresponding news article summaries: 1. Dr. Fauci describes Biden's vaccine mandate as 'moderate' on CNN, emphasizing the need for options for those hesitant to get vaccinated. Republican governors criticize the mandate as an infringement on individual liberty and plan to challenge it in court. The mandate's impact on those previously infected with COVID sparks debate, with Fauci acknowledging the complexity of the issue. 2. A right-leaning article criticizes Dr. Fauci for supporting what they view as President Biden's overly lenient vaccine mandate approach, highlighting Fauci's perceived lack of stringency and labeling Biden's actions as dictatorial, while emphasizing the need to expose the true intentions of political figures like Fauci and Biden. 3. Dr. Fauci characterizes President Biden's vaccine mandate as moderate, emphasizing the option for testing as a compromise for those hesitant to get vaccinated, reflecting a right-leaning perspective on the level of stringency in vaccine mandates.

Table 19: Shows a negative example. The generated partisan summary is incorrect. This is primarily attributed to inconsistent news article summaries (2 and 3), and LLM's failure to identify cited information in the news article.

Points of disagreement	
Left	Right
Title: Biden's Climate Agenda as Economic Focus Description: Climate change has become a central component of President Biden's economic agenda, emphasizing the integration of environmental goals with job creation and economic growth, reflecting a progressive approach to tackling climate issues.	Title: Economic Concerns and Job Losses in Fossil Fuel Industries Description: Right-leaning critics express worries about job losses and economic impacts on industries like oil, natural gas, and coal due to Biden's climate agenda, contrasting it with the potential benefits of transitioning to cleaner energy sources.

Table 20: Shows an example where both the partisan perspectives are in disagreement with each other.

Points of agreement	
Left	Right
Title: Emphasis on U.S. leadership in climate action Description: Buttigieg and Biden stress the importance of U.S. leadership in climate action, contrasting it with China's commitments and urging other nations to follow suit. This narrative positions the U.S. as a moral leader challenging global counterparts to enhance their climate commitments.	Title: Emphasis on U.S. leadership in climate action Description: The focus on the importance of the U.S. leading the way in climate action and challenging other countries to catch up reflects a right-leaning perspective that prioritizes American leadership and influence in global climate initiatives.

Table 21: Shows an example where both the partisan perspectives are in agreement with each other for an event from the issue *Climate Change*.

Prompt to characterize a cluster
Given a set of news article excerpts taken from the same news event, the task is to analyze the articles with the intent to identify a high-level concept that captures the key aspect of discussion related to that event. The concept should be indicative of one of the main discussion angles related to the event, and not very specific to entities mentioned in the articles. The concept should be accompanied by a summary, which should not be a mere concatenation of articles.

Table 22: Prompt to characterize cluster candidate. We prompt the LLM in a two-shot setting.

Prompt to merge two clusters
Given two aspects from the same news event, you need to analyze them with the intent to understand if they are focusing on the same aspects of that event. You should compare the key emphasis of the aspects and their implications to decide if they are the 'same', or 'different'. Refrain from providing any additional explanations other than the label.

Table 23: Prompt to merge clusters. We merge two clusters if their aspects are the same.

F.2 Remove Incoherent Clusters

We note that HDBSCAN algorithm provides us with an initial set of candidate clusters. For each candidate, we use the aspect associated with the cluster label to validate if the top-3 members that are closest to the cluster label in the embedding space are discussing the same high-level concept. We prompt the LLM using the prompt shown in 24 to remove incoherent clusters.

F.3 Talking Point Membership

After obtaining the cluster labels, which characterize the space of possible talking points. We consider each talking point from the set of all the talking points and assign the closest cluster label based on cosine similarity score. If this score is beyond a threshold value of 0.85, we assign the talking point to that cluster label. Otherwise, the it is discarded but retained in the unclustered pool of talking points.

Prompt to remove inconsistent clusters
You are provided with a few news article excerpts and a key aspect of discussion, all of which are from the same news event.
The task is to analyze if all of the provided news article excerpts are discussing the given key aspect.
Respond with 'yes' if the central theme of discussion in each excerpt align in meaning with the key aspect, and 'no' if there is any variance, refraining from offering any additional explanation.

Table 24: Prompt to remove inconsistent clusters.