A Framework of Narrative Media Framing in Political Discourse

Anonymous ACL submission

Abstract

Narrative frames are a powerful way of conceptualizing and communicating complex, controversial ideas. However, while the impact of narrative framing in media has been widely acknowledged, few NLP studies have considered the aspects which make them an effective framing device. In this paper, we show how elements of narrativity in such frames link to fundamental aspects of framing, and present a framework which formalizes and operationalizes such aspects. We annotate and release a data set of news articles in the climate change domain, and perform extensive experiments with LLMs to test their ability to understand narrative frames and their components. Then, we apply the framework in an unsupervised way to discover components of narrative framing in an unrelated domain of COVID-19 crisis, and show how it can generalize across topics to arrive at insights consistent with theoretical narrative framing analysis.

1 Introduction

Narrative framing is a powerful type of media framing that uses elements of narrativity to highlight some aspects of a complex issue and condense it into a simplified "story" that promotes a particular interpretation (Crow and Lawlor, 2016). The elements of storytelling such as representing an issue through its stakeholders and a conflict between them rather than by describing its aspects directly make narrative frames more effective mechanisms of public influence than topic-like generic and issue-specific frames (Daniels and Endfield, 2009; Rodrigo-Alsina, 2019).

The power of narrative framing comes from its ability to draw the reader's attention to more nuanced aspects of an issue and thus instill a very precise interpretation which can be different from the "default" reading inferrable from its generic or issue-specific frame. To give an example, the text in

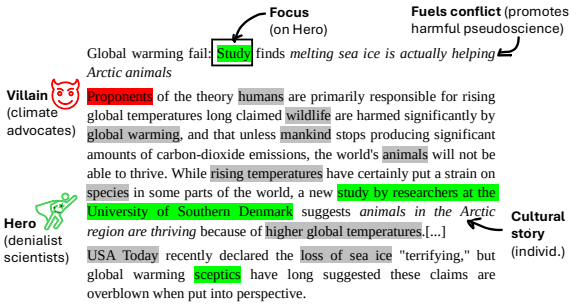

Figure 1: An except from a news article, with `hero` marked in green and `villain` in red. Entities that are not main characters are grayed out. The box shows the focal character (here, `hero`). The phrases in italic are cues which show that the article has an *individualistic* cultural story ("the nature can fix itself") and that it *fuels conflict* by actively promoting bad science.

Figure 1 frames the topic of climate change through a "Polar bear" issue-specific frame (Bushell et al., 2017)[1] which describes the negative effects of climate change on animals (*rising temperatures have certainly put a strain on species*). However, this is not what the text is trying to convince the reader of: it depicts climate scientists and activists as incorrect, while presenting pseudo-scientists from a hero-like angle. Essentially, it uses devices of narrative framing to replace the default interpretation arising from a topic-like frame ("animals are victims of climate change and humans are villains") with an opposing idea that animals are doing fine and people who think otherwise are villains.

While the importance of framing narratives for the communication and perception of news has been widely recognized in the social sciences (Shanahan et al., 2011), automatic framing analysis still mostly conceptualizes frames in a topic-like fashion (Ali and Hassan, 2022; Ot-

[1]None of the widely adopted generic frames such as by (Boydstun et al., 2013) or (Semetko and Valkenburg, 2000) is readily applicable to this text, which shows their interpetative limitations.

makhova et al., 2024). What distinguishes a frame from a topic? For framing to occur, an issue must be ambivalent (have alternative interpretations) for it to be frame-able (Sniderman and Theriault, 2004). A frame must evoke a larger interpretative context (schemata) that goes beyond information inferrable from the text (Scheufele and Scheufele, 2010). The topical component of a frame relates to Entman (1993)'s frame component of the "problem statement". But only in combination with Entman's other components (moral evaluation, conflict definition and prescribed "treatment", or resolution) a message becomes a frame. We present a formalization of narrative framing that comprises all aspects.

We do so by drawing upon insights from narratology research, and from social and media studies that focus on narrative framing, to establish a framework which allows to identify narrative frames and to distinguish similar narrative frames from each other. In particular, we show how issue ambiguity arises from character role assignment (Hero, Villain, Victim) to actors in the narrative frame, and from focusing on one of such characters. For example, in the article in Figure 1 climate activists receive a role of a Villain, while denialist scientists are regarded as Heros, and the narrative focuses on the heroes' "contributions". Next, we map the relationships between characters to a general definition of conflict and resolution (the article in Figure 1 fuels the climate crisis rather that its resolution). Finally, as framing can occur only when the frame evokes a wider set of associations and believes already existing in the receiver's perception (Nelson et al., 1997), narrative frames are linked to established "cultural stories" that define the attitude towards external control and the sense of unity with the group (Thompson, 2018). The text above is an example of individualistic cultural story, which implies that no control is needed and the society does not need to act as one.

We apply this framework to analyze media framing of two distinct public issues – climate change and COVID-19. In particular, we make the following contributions:

1. We define elements of narrativity that are essential for narrative framing and are aligned with the definition of a "frame".

2. We show that our framework composed of such elements can be applied to different domains: climate change and COVID 19.

3. We show how our framework enables to annotate narrative frames in the news reliably and efficiently, and to perform exploratory framing analysis.

4. We release a corpus of articles about climate change annotated with narrative frames, and use it to evaluate LLMs on their ability to automatically predict the components of our framework.

2 Background

Narratives in political communication Following Fisher (1984)'s seminal paper coining the term 'homo narrans' to illustrate the importance of storytelling for society, narratives in political communication have attracted substantial research attention (see also Bennett and Edelman (1985); Patterson and Monroe (1998)), exposing its effects from a critical vehicle in deliberative democracy (Boswell, 2013) to its use persuasive device (Skrynnikova et al., 2017). Similar to related concepts such as 'framing', a principled and empirically testable definition of 'narrative' has long been lacking. However, recent work has progressed in developing frameworks that are testable and amenable to computational modelling (Shenhav, 2005; Robert and Shenhav, 2014), most prominently the Narrative Policy Framework (NPF; Jones et al. (2023)), which we build on in this work. The NPF has been particularly instrumental in studying climate change narratives (Fløttum and Gjerstad, 2017), and identifying dominant narratives in the discourse (Bushell et al., 2017; Bevan, 2020). This paper refines the NPF into a structured framework suitable for automatic prediction.

Narrative framing and NLP Narrative framing intersects the concepts of storytelling and framing, i.e, the presentation of information in a way to evoke a specific association in the audience. Automatic narrative understanding has attracted substantial attention in NLP (Piper et al., 2021), however, has focussed mostly on fictional narratives (Bamman et al., 2014; Iyyer et al., 2016), on personal narratives in social media (Lukin et al., 2016; Shen et al., 2023) or specific elements such as event chains (Chambers and Jurafsky, 2009). Few works have considered the intersection of stories and framing (Levi et al., 2022). While narrative framing research in the social sciences is strongly grounded in the NPF, it is yet to gain recognition and adoption in NLP approaches. Closest to our work are

Stammbach et al. (2022) and Frermann et al. (2023) who study some narrative elements of framing devices (such as entities framed as heroes or victims), but do not model full narrative frames. In addition to entities Gehring and Grigoletto (2023) model relationships between them such as "harm" or "protect"; however, their approach is closer to named entity and relation extraction than to narrative frame understanding, and does not map the specific identified elements to more abstract ideas such as narrative frames.

3 Components of narrative framing

We motivate our three core components which define a narrative frame. Each component contributes to the framing mechanism, by resolving the ambivalence by assigning moral evaluation to stakeholders (Characters), capturing the conflict and resolution aspect of a frame (Conflict and resolution), and evoking a wider set of cognitive schemata and cultural associations (Cultural stories).

3.1 Characters

Characters and their prototypical roles have been studied extensively in narratology (starting from formalist and structuralist approaches such as Propp (1968) and Greimas (1987)), and were adopted as a simplified `hero`, `villain`, and `victim` (HVV) triad by social sciences as part of Narrative Policy Framework (NFP) (Shanahan et al., 2018)[2]. In particular, NFP prescribes that a narrative frame should include **at least one prototypical character**, i.e. one or more of HVV role slots should be filled by a prominent entity. By assigning an entity to a particular role, we resolve the issue ambivalence by conveying our moral judgement of that entity, as required by Entman's definition of frame (Entman, 1993). In particular, the reader's interpretation of the article depends on whether a particular entity (say, *climate advocates* as in Figure 1) is framed as a `hero` (their actions are evaluated as beneficial), a `villain` (as in our example), or `victim` (of criticism or attacks by denialists).

Though a text can have multiple candidates for each of HVV roles, we follow narratology approaches that focus on distinguishing between main characters and other entities (Jahan and Finlayson, 2019), and use the *single most central character* fulfilling the respective role to represent a narrative frame. Figure 1 illustrates this, where the main characters are highlighted in color, while less central entities are grayed out. Moreover, to be able to compare instances of a particular narrative frame across texts with different people and events, we abstract away from specific characters and instead use stakeholders (common people, elites, etc.) they represent. The taxonomy of such stakeholders can either be inherited from theoretical literature (as we do in Section 4.1) or derived automatically from texts (as in Section 4.2).

To fully differentiate narratives, in addition to assigning characters to roles it is necessary to identify the **focus** on either `hero`, `villain`, or `victim`, which results in "heroic", "blaming", and "victimizing" narrative frames, respectively. For example, two distinct narrative frames can both have *climate activists* as a `hero` and *government* as `villain`, but focus either on criticizing the government ("blaming") or praising the efforts of activists in opposing it ("heroic") (examples from Bevan (2020)).

3.2 Conflict and resolution

Conflict/resolution[3] is a central element of a narrative frame, tying it to Entman's (Entman, 1993) criteria of framing which state that a frame should define a conflict, contain its moral evaluation, point to its cause, or prescribe a solution. In particular, `hero`, `villain` and `victim` can either try to exacerbate the issue or resolve it, or assign the causes of a problem to someone else and give their moral judgement. Accordingly, we conceptualize conflict and resolution as a four-way distinction: the characters in a narrative frame can either *fuel conflict* (perform actions that exacerbate the issue), *fuel resolution* (perform actions that help to resolve the issue), *prevent conflict* (oppose actions that exacerbate the issue), or *prevent resolution* (oppose actions that help to resolve the issue).

In NLP, relations between characters have a long history of research (Agarwal and Rambow, 2010; Shahsavari et al., 2020), including studies which specifically looks at conflicts (Han et al., 2019; Olsson et al., 2020). In comparison to them, our framework provides a more abstract way of repre-

[2] In NLP, character (or "agent") identification has attracted substantial attention, both from the narratology side (see (Piper et al., 2021) for a review) and, less extensively, from NFP side (Frermann et al., 2023).

[3] Here we understand conflict as a problem or an issue which characters strive to either escalate or resolve (in opposition to each other), rather than a driving force of a plot (Prince, 2003) or breaking point in its canonicity (Bruner, 1991), as understood in narratology.

senting the conflict expressed in a narrative frame, which combines the attitude towards the issue (pro-conflict vs pro-resolution) with the level of intentionality and direct expression of that attitude (i.e. actively perform actions that support one's side, or oppose the actions of the other side). Such definition of conflict/resolution based on abstract categories rather than on specific actions or events makes our approach generalizable across topics, as we show in Sections 4.1 and 4.2.

3.3 Cultural stories

Frames are distinguished from "unframed" types of communication by their ability to evoke a wider set of concepts, associations and judgements which already exist in the audience's perception (Scheufele and Scheufele, 2010). Narrative frames do this by mapping a particular combination of characters and conflict/resolution to one of four larger schemata of interpretation, which in social studies are referred to as **cultural value stories** (Thompson, 2018). Cultural stories define to what degree our actions are controlled by external factors and by the sense of belonging to a particular group (Douglas, 2007).[4] Depending on the combination of these two factors, a narrative frame can be *fatalist* (where people are at the mercy of forces outside their control, such as natural disasters or fate), *hierarchical* (where people are bound by social prescriptions and external control, such as government), *individualistic* (where social ties are loose and people reject the necessity of external control), or *egalitarian* (where people take collective action, opposing external control) (Figure 2). The effects of cultural stories on public behavior (and thus their framing power) are substantial: as an example, individualist and egalitarian stories lead to worse survival outcomes than a hierarchical story in life-critical situations such as onset on COVID-19 pandemic (Güss and Tuason, 2021).

To the best of our knowledge, cultural stories, or more generally schemas aiding interpretation, have not been explored in NLP. However, many NLP studies (Finlayson, 2012; Tangherlini et al., 2020) draw upon related concepts of narrative archetypes and schema as overarching, culturally repetitive plots or narrative elements Frye (1957); Propp (1968). In contrast, we focus on framing and its link to a well-defined space of cultural values which have been shown to affect perception and behavior.

[4]Thus, all narrative frames are "stories", i.e. contain some elements of narrativity such as characters and plot, reduced to conflict and resolution. However, not all potential "stories" can be used as narrative frames: in order to be such, they need to map to a broader, pre-existing context dictated by a cultural story.

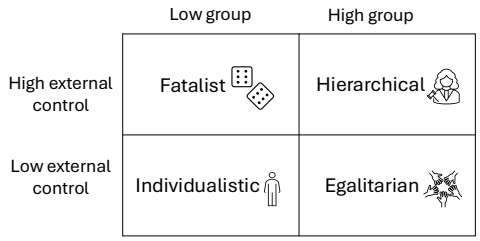

Figure 2: Cultural stories across dimensions of **external control** (grid) and belonging to a **group**

4 Narrative frame analysis

In the sections below we apply our framework to perform narrative framing analysis on two topics: climate change and COVID-19. First, in Section 4.1 we use it to annotate structures of narrative frames in news articles, and then map them to well-known narrative frames for climate change domain. We use this corpus to evaluate the ability of multiple LLMs to understand narrative frame components and the frames themselves. Then, in Section 4.2 we show how the framework can be generalised to a domain without an established repertoire of narrative frames, where the goal is to *discover* frames rather than classify them. Specifically, we present an exploratory analysis of politician's speeches during the onset of COVID-19 and show how our framework surfaces insights consistent with prior analyses.

4.1 Supervised approach: climate domain

4.1.1 Dataset

We manually annotate 100 articles randomly selected from a dataset of news stories on the topic of climate change (Frermann et al., 2023). Below we outline the steps of the annotation process, using the example in Figure 1.

First, we identify candidate entities for the `hero`, `villain`, and `victim` roles, and select at most one main character per role (as described in Appendix B). In our example, since `hero`, `villain`, and `victim` should align with the article's perspective, we remove potential victims like *animals* since the author believes they actually benefit from higher temperatures. Then, using a previously established taxonomy of stakeholder types for the climate change domain (Frermann et al. (2023); details in Appendix A), we map the text spans that

represent characters to labels representing general classes of actors. Thus, we arrive at *science experts* (from the skeptics side) as `hero` and *environmental activists* as `villain`. To determine the focus, we rely on rhetoric devices and discourse structure of newspaper articles, namely the inverted pyramid where the most important content is usually placed at the beginning, and the relative proportion of text devoted to the different roles. Since the title highlights the research of climate skeptics, and much of the article's content is devoted to describing it, we determine that the `focus` is on the `hero`. Next, since the article explicitly promotes bad science harmful for climate (rather than only criticising actions of climate activists), it fuels `conflict`. Finally, as the article implies that nature is resilient and no actions are necessary, it corresponds to *individualistic* cultural story.

We apply this framework to news articles to annotate the structure of their narrative frames, and use the same framework to determine the components of known narrative frames in the climate change literature (Fløttum and Gjerstad, 2017; Bushell et al., 2017; Bevan, 2020).Then, we map the article structures to the structures of known narrative frames to arrive at the final narrative frame label for the article. This process resulted in defining 16 structurally distinct narrative frames, which are described in detail in Appendix D. Overall, this structure points to a denialist narrative frame "No need to act". Full dataset statistics are provided in Appendix E.

We perform annotation following a rigorous process and achieve reliable inter-annotator agreement on all elements of the framework, as described in in Appendix C. We find that our framework improves understanding and recognition of narrative frames by annotators, as well as reduces the annotation time Appendix C.2. As we discovered the narrative frames in an inductive, bottom-up way, we cannot claim that our dataset contains a complete set of narrative frames in the climate discourse. However, we note that it covers the majority of narrative frames mentioned in social studies literature, including frames very similar in content but differing in structure, and thus can be used as a representative dataset for testing narrative frame detection in this domain.

4.1.2 Task definition

We use the resulting dataset of 100 articles to test narrative frame understanding capabilities of LLMs. We consider the following prediction tasks, given the full article text as input:

- Predicting the stakeholder category for `hero`, `villain`, and `victim` separately for each of the characters as one of 10 stakeholder classes (*government*, *climate activists*, etc; see Section 3.1). To choose a stakeholder for a character correctly, a model needs to perform several reasoning steps: determining if an entity is framed as a `hero`, `villain`, or `victim`, aggregating mentions of entities across the text to determine which of potential candidates is a main `hero`, `villain`, or `victim`; and finally determining to which stakeholder category this character belongs.

- Predicting the `focus` out of 3 classes (`hero`, `villain`, or `victim`). This task tests if a model is able to determine if the narrative frame is "heroic", "blaming", or victim-centered.

- Predicting `conflict` out of 4 classes (*fuel conflict, fuel resolution, prevent conflict, prevent resolution*, see Section 3.2). The model must understand the general intent of the narrative frame (if it pushes towards resolution of the crisis, or exacerbates it), and if it does it by actively supporting one side or by criticising the opposite side.

- Choosing a `cultural story` out of 3 classes (*individualistic, egalitarian, or hierarchical*, see Section 3.3)[5]. To do so, the model should understand if the text implies collective vs individualistic action, and approval or disapproval of external control (such as from the government).

- Choosing one of 16 `narrative frames` based on their short descriptions sources from social studies literature (full list in Appendix D).

4.1.3 Models and prompts

We use our tasks to test narrative frame understanding of 6 LLMs of different size and complexity: GPT4o, GPT4o1, Mixtral, Llama, Gemini and Claude Sonnet.[6] We set temperature=0 (except for GPTo1 which does not allow to control generation) in all experiments to ensure deterministic

[5]Though Section 3.3 introduces 4 cultural stories (Jones, 2014), the fatalist story is not present in our data set so we exclude it from experiments for fair evaluation.

[6]Versions used: gpt-4o-2024-11-20, o1-preview-2024-09-12, Mixtral-8x7B-Instruct-v0.1, Llama-3.1-8B-Instruct, gemini-1.5-flash, Sonnet 3.5. Model sizes are provided in Appendix F.

	Hero (10)	Villain (10)	Victim (10)	Focus (3)	Conflict (4)	Story (4)	Narrative (16)
Baseline	0.079	0.08	0.135	0.231	0.135	0.19	0.021
GPT4o	0.325	0.454	0.266	0.656	0.332	0.574	0.258
GPT4o1	0.363**	0.527*	0.455*	0.718**	**0.549**	**0.595**	0.330*
Mixtral	0.237	0.073	0.257	0.402	0.353	0.431	0.171
Llama	0.271	0.156	0.336	0.568	0.379	0.449	0.181
Gemini	0.326	0.292	0.230	0.635	0.361	0.482	0.319
Sonnet	**0.353**	**0.530**	**0.469**	**0.688**	0.399	0.561	**0.339**

Table 1: Zero-shot performance of 6 models in terms of macro- averaged F1 across 7 narrative understanding tasks. The number in brackets after the task's name indicated the number of classes in it. The baseline is calculated by using the most frequent label for a particular task as a predicted class. Results that had high (over 0.02) or very high (over 0.05) standard deviation across 5 runs are marked with * and ** respectively. The best performing models (considering variance) are in bold.

outputs. We perform each experiment for 5 runs to ensure there is no substantial variance in the results. With the exception of GPT4o1, which shows high variance on most tasks, all models have zero (Mixtral, Llama) or near-zero (GPT4o, Gemini, Sonnet) variance across runs, which allows to compare the averaged results between models.[7]

The prompts used for each of the tasks are listed in Appendix G. The text of the prompts is based on descriptions of particular classes (stakeholders, culture stories, narrative frames etc.) in the social studies literature. Prompts for HVV characters are domain-specific, i.e. they are based on a list of entities important for the climate change domain (we show how to generalize this approach by creating such list automatically in Section 4.2). Conversely, prompts for Focus, Conflict, and Cultural story tasks are domain-agnostic, describe the classes in general terms (e.g., *INDIVIDUALISTIC: this story assumes that the situation cannot be controlled externally, and no group actions are necessary*). We use the most abstract prompts possible to ensure the approach is generalizable, but we also found that abstract prompts lead to better performance compared to prompts specifically describing how a particular conflict or cultural story is manifested in the climate change debate.

4.1.4 Results

Results in Table 1 show that no single model consistently performed best (or worst) across all tasks. Mixral and Llama are the weakest, especially in stakeholder prediction for for `hero` and `villain` where both models overpredict entities that are stereotypical heros and villains for this topic. For instance, they select "environmental activists" as heroes and "pollution" as villain, despite the fact that they rarely occur in these roles in our articles. In a similar way, they overpredict rare narrative frames as "Carbon fuelled expansion" which claims that fossil fuels are necessary for economy, probably by matching them on topical vocabulary such as "fossil fuels". Thus, weaker models tend to overgeneralize their "knowledge" about the topic, disregarding the content and intent of articles.

The strongest models, Sonnet and GPT4o1, tie in terms of results, but Sonnet is faster (8 min vs 26 min on average for 100 articles), more stable in terms of variance, and less costly (due to a large number of internal chain-of-thought tokens generated by GPT4o1). However, GPT4o1 excels in tasks which require understanding the overall "gist" of the text, such as predicting Cultural story and Conflict. Gemini and GPT4o lie in the middle in terms of performance, with GPT4o doing better in terms of HVV stakeholders and Cultural Story, but substantially worse in terms of Narrative classification, where it significantly overpredicts denialists narrative frames. That problem, however, is not specific to GPT4o, with all models performing poorly on detecting narrative frames based on their description and tending to excessively predict one or two classes.

We perform experiments to optimize the prompt and help models learn from examples (see Appendix I), but they do not lead to performance gains, which shows the difficulty of the tasks.

The effects of number of classes The difficulty of the Narrative task is confounded by the number of classes that need to be distinguished (16). To test if the performance would increase if the model is asked to choose between a small number of classes,

[7]When comparing with gpt-4o-2024-11-20, we used the worst results of the models rather than average to account for large variance.

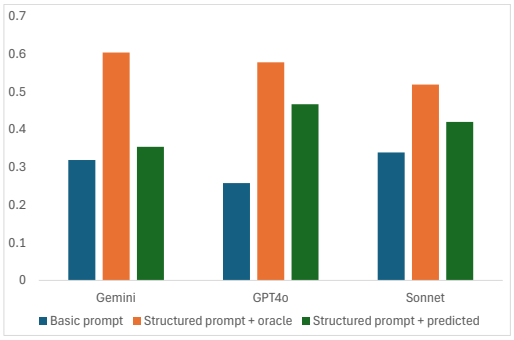

Figure 3: Predicting narrative frames using oracle (human-annotated) and noisy (predicted) labels for their `hero`, `villain`, `victim`, and `focus`; the results are macro-averaged F1.

we select a sample of three frequent, but similar narratives - "12 Years to save the planet", "We are all going to die", and "Gore" (see Appendix D), and modify the Narrative prompt to include descriptions only of these three classes. However, this increases performance only minimally (from F1 of 0.258 to 0.27 for GPT-4o) and nowhere near the level for tasks with a comparable number of classes such as Focus and Conflict. Moreover, even for this simplified task there is a tendency to predict one class, and one of the classes is never chosen correctly (Figure 8 in Appendix).

4.1.5 Predicting narrative frames with labels for their structure

In this section, we explore if using narrative frame structures (such as specifying `hero`, `villain`, `victim`, and `focus`) can improve narrative frame classification. For these experiments we use three models (the strongest Sonnet and middle-grade GPT-4o and Gemini[8]) in zero-shot mode. For each of the narrative frames definitions we add an informal description of typical stakeholders for `hero`, `villain`, `victim` (as listed in Appendix D and who the frame is focusing on (`focus`) (see Appendix H).

Next, for each input article we add labels that denote `hero`, `villain`, and `victim` stakeholder category, as well as *focus*, to explicitly represent the structure of the narrative frame. With *oracle* (manually-annotated) labels the performance improves substantially across models (most notably in GPT4o), which shows that explicit structure is a more reliable cue for predicting the narrative than its descriptor. We repeat this experiment with noisy labels predicted by the model with the best zero-shot prompting results (Sonnet), and observe gains compared to the prompt without structure. The models behavior, however, is quite different: Gemini struggles to gain useful information from noisy labels, Sonnet benefits least of all from the structure, but does accommodate for the noisiness, while GPT4o, despite being the weakest on this task, benefits most from structure and noisy labels, achieving substantially better performance than the best-performing model (Sonnet)[9].

To examine how introduction of structure affects the narrative frame prediction, and analize which narrative frames are hard for models to predict even when they are given correct labels for main characters, we compare confusion matrices for Narrative classification with a basic prompt, and with a structured prompt and oracle character labels (Appendix J). We observe that before the introduction of structure the predictions are quite sporadic and scattered across the matrix; i.e. both the predictions and errors are not systematics. With the structure, however, we see clear patterns of consolidation: first, some narratives frames which have an unique structure (such as "Official declare emergency" which, unlike most others, frames government as a `hero`) are now predicted perfectly or near-perfectly rather than randomly as before. Errors, too, occur due to confusion of a handful of structurally similar narrative frames. Most of the errors are due to the confusion of two frames that both focus on criticizing the government (`villain`), but have different cultural stories: one of them ("12 years to save the Earth") calls for even more governmental control (*hierarchical*), while the other ("All talk no action") opposes the government (*egalitarian*). This highlights the importance of other aspects in addition to characters to differentiate between narrative frames, and opens the possibility of further improvement by incorporating them.

4.2 Unsupervised approach: COVID-19 narrative frames

In this section we apply our narrative frame structures taxonomy to texts with a different topic and style – politicians addresses regarding COVID-19 – to demonstrate its generalizability to other domains. We show how models and prompts devel-

[8]We exclude GPT-4o1 due to its high costs and instability.

[9]These gains, however, occur when the labels are accurate enough: we observed only minimal gains or drops in performance when using predictions from less strong models.

oped within the supervised approach can be applied to analyze narrative frame components in an unsupervised way.

4.2.1 Dataset and model

We collect transcripts of head-of-state addresses regarding the onset of COVID-19 dating from February to end of July 2020, for three countries: Germany (Angela Merkel; N=12), UK (Boris Johnston; N=24) and Australia (Scott Morrisson; N=6).[10]

We use the most reliable model identified in Section 4.1 (Claude Sonnet 3.5) in a zero-shot scenario. Since the prompts for `focus`, `conflict`, and culture story developed in Section 4.1 are domain agnostic, we apply them without changes, only substituting the topic name for "COVID 19". However, since the set of stakeholders is likely to be different for this topic, we modify the HVV prompts by replacing the classes with a list of topic-specific stakeholders. We compile this list automatically by generating them from the speeches: first we ask the LLM to extract and merge entities which are likely to represent `hero`, `villain`, and `victim`, then combine the extracted candidates from all speeches and cluster them into groups (prompts in Appendix K). We arrive at a set of 8 stakeholders, some of which are generic and shared with the climate change domain (*government, general public*), while the majority are unique and topic-specific (*vulnerable population, healthcare*, etc). The final set of stakeholders corresponds to prominent stakeholders identified as `hero`, `villain`, `victim` in studies on narrative framing in these speeches (Bernard et al., 2021; Mintrom et al., 2021).

4.2.2 Narrative frame analysis

We apply our approach to discover differences and commonalities in framing of politicians speeches regarding COVID-19.

First, all speeches across all three politicians were identified as *hero-focused* and *promoting resolution*, which is not surprising given the fact that they are all mobilizing narratives that suggest specific actions to solve the crisis and praise the role of heroes. Similarly, the `villain` is consistently detected as "pandemics", and `victim` is "general public", especially "vulnerable populations", and, later in the period, "economy".

However, the stakeholders that are pinpointed as `hero` are different for these three politicians: while all of them recognize the role of "healthcare workers", Merkel's speeches also highlight the role of "general public", and, later in the pandemic, of "global efforts". On the other hand, Morrison's speeches heavily revolve around the role of "government" as a hero, as well as mention "science experts". Such distinction corresponds to theoretical analyses of these speeches, which assert that chancellor Merkel's recognized the value of combined efforts of German public, as well as global efforts (Mintrom et al., 2021), while prime minister Morrison often used reassuring framing relying on the role of science in pandemic management and using the imagery of Australia as "lucky country" (Bernard et al., 2021). Similarly, the analysis of predicted cultural stories reveals that Morrison predominantly used *hierarchical* cultural stories (*'Government and following social prescriptions plays the biggest role in managing the crisis'*), Merkel had a larger proportion of *egalitarian* narrative frames than others (*'We must act as one to combat the crisis'*), while Johnson was the only one who alluded to *individualistic* cultural story (*'Take care of yourself and your family'*). Again, these insights align with previous theoretical analyses (Mintrom et al., 2021).

In sum, we showed that LLMs can (to an extent) identify core components of narrative frames when evaluated against human-labelled data, and that our framework generalizes across topics and to scenarios of discovering frame structures in a previously unseen domain.

5 Conclusion

We presented a rigorous formalization and taxonomy of components of narrative framing. Our method allows to inductively detect narratives in news articles in terms of their character types, focus, conflict, and underlying cultural story. A high-quality data set of 100 labeled articles serves as a benchmark and basis for future annotation bootstrapping. We showed that our framework results in promising performance of automatic narrative prediction with LLMs, laying a foundation for the important research agenda of large-scale studies of the manifestation and effects of narrative frames. Moreover, we showed that our framework is generalizable to other topics and can assist in exploratory framing analysis without requiring a labeled dataset.

[10] Sources: https://www.bundesregierung.de/breg-en/service/archive/, https://www.gov.uk/government/speeches/, https://www.pm.gov.au/media.

6 Limitations

We acknowledge the small size of our data set relative to NLP benchmarks, but emphasize the difficulty of annotating news articles at this level. We prioritize depth over breadth, and our data set can serve both as a benchmark and a high-quality starting point for bootstrapping other story annotations.

Because our approach is inductive / bottom-up we cannot guarantee that the narratives we found cover all possible active narratives or reflect the true narrative distribution. However, since our inductive narratives overlapped with a large part of narratives described in the literature, we are confident that they are representative and comprehensive.

Additional LLM experiments, with larger example pools or advanced reasoning techniques may lead to further improvements but are outside the scope of this work. We showed that incorporating narrative structure into prompts improves performance more substantially than models with advanced reasoning abilities. Future work, however, may want to combine it with such models and techniques.

References

Apoorv Agarwal and Owen Rambow. 2010. Automatic detection and classification of social events. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1024–1034.

Mohammad Ali and Naeemul Hassan. 2022. A survey of computational framing analysis approaches. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9335–9348.

David Bamman, Ted Underwood, and Noah A. Smith. 2014. A Bayesian mixed effects model of literary character. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 370–379, Baltimore, Maryland. Association for Computational Linguistics.

W Lance Bennett and Murray Edelman. 1985. Toward a new political narrative. *Journal of communication*.

Natalie Reyes Bernard, Abdul Basit, Ernesta Sofija, Hai Phung, Jessica Lee, Shannon Rutherford, Bernadette Sebar, Neil Harris, Dung Phung, and Nicola Wiseman. 2021. Analysis of crisis communication by the prime minister of australia during the covid-19 pandemic. *International Journal of Disaster Risk Reduction*, 62:102375.

Luke D Bevan. 2020. Climate change strategic narratives in the united kingdom: emergency, extinction, effectiveness. *Energy research & social science*, 69:101580.

John Boswell. 2013. Why and how narrative matters in deliberative systems. *Political studies*, 61(3):620–636.

Amber E Boydstun, Justin H Gross, Philip Resnik, and Noah A Smith. 2013. Identifying media frames and frame dynamics within and across policy issues.

Jerome Bruner. 1991. The narrative construction of reality. *Critical inquiry*, 18(1):1–21.

Simon Bushell, Géraldine Satre Buisson, Mark Workman, and Thomas Colley. 2017. Strategic narratives in climate change: Towards a unifying narrative to address the action gap on climate change. *Energy Research & Social Science*, 28:39–49.

Nathanael Chambers and Dan Jurafsky. 2009. Unsupervised learning of narrative schemas and their participants. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 602–610, Suntec, Singapore. Association for Computational Linguistics.

Deserai A Crow and Andrea Lawlor. 2016. Media in the policy process: Using framing and narratives to understand policy influences. *Review of Policy Research*, 33(5):472–491.

Stephen Daniels and Georgina H Endfield. 2009. Narratives of climate change: introduction. *Journal of Historical Geography*, 35(2):215–222.

Mary Douglas. 2007. A history of grid and group cultural theory. *Toronto, Canada: University of Toronto*.

Robert M Entman. 1993. Framing: Toward clarification of a fractured paradigm. *Journal of communication*, 43(4):51–58.

Mark Mark Alan Finlayson. 2012. *Learning narrative structure from annotated folktales*. Ph.D. thesis, Massachusetts Institute of Technology.

Walter R Fisher. 1984. Narration as a human communication paradigm: The case of public moral argument. *Communications Monographs*, 51(1):1–22.

Kjersti Fløttum and Øyvind Gjerstad. 2017. Narratives in climate change discourse. *Wiley Interdisciplinary Reviews: Climate Change*, 8(1):e429.

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)*, pages 8712–8732. Association for Computational Linguistics.

Northrop H Frye. 1957. *Anatomy of criticism*. Princeton.

Kai Sebastian Gehring and Matteo Grigoletto. 2023. Analyzing climate change policy narratives with the character-role narrative framework. *SSRN Electronic Journal*, 10429.

Algirdas Julien Greimas. 1987. On meaning: Selected writings in semiotic theory. *Actants, Actors, and Figures*.

C Dominik Güss and Ma Teresa Tuason. 2021. Individualism and egalitarianism can kill: how cultural values predict coronavirus deaths across the globe. *Frontiers in Psychology*, 12:620490.

Xiaochuang Han, Eunsol Choi, and Chenhao Tan. 2019. No permanent Friends or enemies: Tracking relationships between nations from news. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1660–1676, Minneapolis, Minnesota. Association for Computational Linguistics.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. LoRA: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.

Mohit Iyyer, Anupam Guha, Snigdha Chaturvedi, Jordan Boyd-Graber, and Hal Daumé III. 2016. Feuding families and former Friends: Unsupervised learning for dynamic fictional relationships. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1534–1544, San Diego, California. Association for Computational Linguistics.

Labiba Jahan and Mark Finlayson. 2019. Character identification refined: A proposal. In *Proceedings of the First Workshop on Narrative Understanding*, pages 12–18, Minneapolis, Minnesota. Association for Computational Linguistics.

Michael D Jones. 2014. Cultural characters and climate change: How heroes shape our perception of climate science. *Social Science Quarterly*, 95(1):1–39.

Michael D Jones, Aaron Smith-Walter, Mark K McBeth, and Elizabeth A Shanahan. 2023. The narrative policy framework. In *Theories of the policy process*, pages 161–195. Routledge.

Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan, Saiful Haq, Ashutosh Sharma, Thomas T Joshi, Hanna Moazam, et al. 2023. Dspy: Compiling declarative language model calls into self-improving pipelines. *arXiv preprint arXiv:2310.03714*.

William F Lamb, Giulio Mattioli, Sebastian Levi, J Timmons Roberts, Stuart Capstick, Felix Creutzig, Jan C Minx, Finn Müller-Hansen, Trevor Culhane, and Julia K Steinberger. 2020. Discourses of climate delay. *Global Sustainability*, 3:e17.

Effi Levi, Guy Mor, Tamir Sheafer, and Shaul Shenhav. 2022. Detecting narrative elements in informational text. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1755–1765, Seattle, United States. Association for Computational Linguistics.

Stephanie Lukin, Kevin Bowden, Casey Barackman, and Marilyn Walker. 2016. PersonaBank: A corpus of personal narratives and their story intention graphs. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1026–1033, Portorož, Slovenia. European Language Resources Association (ELRA).

Michael Mintrom, Maria Rost Rublee, Matteo Bonotti, and Steven T Zech. 2021. Policy narratives, localisation, and public justification: responses to covid-19. *Journal of European Public Policy*, 28(8):1219–1237.

Thomas E Nelson, Zoe M Oxley, and Rosalee A Clawson. 1997. Toward a psychology of framing effects. *Political behavior*, 19:221–246.

Fredrik Olsson, Magnus Sahlgren, Fehmi ben Abdesslem, Ariel Ekgren, and Kristine Eck. 2020. Text categorization for conflict event annotation. In *Proceedings of the Workshop on Automated Extraction of Socio-political Events from News 2020*, pages 19–25, Marseille, France. European Language Resources Association (ELRA).

Yulia Otmakhova, Shima Khanehzar, and Lea Frermann. 2024. Media framing: A typology and survey of computational approaches across disciplines. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15407–15428, Bangkok, Thailand. Association for Computational Linguistics.

Molly Patterson and Kristen Renwick Monroe. 1998. Narrative in political science. *Annual review of political science*, 1(1):315–331.

Andrew Piper, Richard Jean So, and David Bamman. 2021. Narrative theory for computational narrative understanding. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 298–311, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Gerald Prince. 2003. A dictionary of narratology. *U of Nebraska P*.

Vladimir Propp. 1968. Morphology of the folktale. *U of Texas P*.

Dominique Robert and Shaul Shenhav. 2014. Fundamental assumptions in narrative analysis: Mapping the field. *The qualitative report*, 19(38):1–17.

Miquel Rodrigo-Alsina. 2019. Talking about climate change: the power of narratives. In *Climate Change Denial and Public Relations*, pages 103–120. Routledge.

Bertram T Scheufele and Dietram A Scheufele. 2010. Of spreading activation, applicability, and schemas: Conceptual distinctions and their operational implications for measuring frames and framing effects. In *Doing News Framing Analysis*, pages 126–150. Routledge.

Holli A Semetko and Patti M Valkenburg. 2000. Framing european politics: A content analysis of press and television news. *Journal of communication*, 50(2):93–109.

Shadi Shahsavari, Pavan Holur, Tianyi Wang, Timothy R Tangherlini, and Vwani Roychowdhury. 2020. Conspiracy in the time of corona: automatic detection of emerging covid-19 conspiracy theories in social media and the news. *Journal of computational social science*, 3(2):279–317.

Elizabeth A Shanahan, Michael D Jones, and Mark K McBeth. 2018. How to conduct a narrative policy framework study. *The Social Science Journal*, 55(3):332–345.

Elizabeth A Shanahan, Mark K McBeth, and Paul L Hathaway. 2011. Narrative policy framework: The influence of media policy narratives on public opinion. *Politics & Policy*, 39(3):373–400.

Mike Shanahan. 2007. Talking about a revolution: climate change and the media.

Jocelyn Shen, Maarten Sap, Pedro Colon-Hernandez, Hae Park, and Cynthia Breazeal. 2023. Modeling empathic similarity in personal narratives. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6237–6252, Singapore. Association for Computational Linguistics.

Shaul R Shenhav. 2005. Thin and thick narrative analysis: On the question of defining and analyzing political narratives. *Narrative inquiry*, 15(1):75–99.

IV Skrynnikova, TN Astafurova, and NA Sytina. 2017. Power of metaphor: cultural narratives in political persuasion. In *7th International Scientific and Practical Conference" Current issues of linguistics and didactics: The interdisciplinary approach in humanities"(CILDIAH 2017)*, pages 279–284. Atlantis Press.

Paul M Sniderman and Sean M Theriault. 2004. The structure of political argument and the logic of issue framing. *Studies in public opinion: Attitudes, nonattitudes, measurement error, and change*, 3(03):133–65.

Dominik Stammbach, Maria Antoniak, and Elliott Ash. 2022. Heroes, villains, and victims, and GPT-3: Automated extraction of character roles without training data. In *Proceedings of the 4th Workshop of Narrative Understanding (WNU2022)*, pages 47–56. Association for Computational Linguistics.

Timothy R Tangherlini, Shadi Shahsavari, Behnam Shahbazi, Ehsan Ebrahimzadeh, and Vwani Roychowdhury. 2020. An automated pipeline for the discovery of conspiracy and conspiracy theory narrative frameworks: Bridgegate, pizzagate and storytelling on the web. *PloS one*, 15(6):e0233879.

Michael Thompson. 2018. *Cultural theory*. Routledge.

A Stakeholder types

We use the following 10 stakeholder types borrowed from (Frermann et al., 2023):
GOVERNMENTS_POLITICIANS: governments and political organizations
INDUSTRY_EMISSIONS: industries, businesses, and the pollution created by them
LEGISLATION_POLICIES: policies and legislation responses
GENERAL_PUBLIC: general public, individuals, and society, including their wellbeing, status quo and economy
ANIMALS_NATURE_ENVIRONMENT: nature and environment in general or specific species
ENV.ORGS_ACTIVISTS: climate activists and organizations
SCIENCE_EXPERTS_SCI.REPORTS: scientists and scientific reports/research
CLIMATE_CHANGE: climate change as a process or consequence
GREEN_TECHNOLOGY_INNOVATION: innovative and green technologies
MEDIA_JOURNALISTS: media and journalists

B Deriving main characters

As we explain in Section 3.1, not all entities in an article represent its main `hero`, `villain`, or `victim`. To be able to reliably and consistently identify the main characters, we adhere to the following process:

(1) We consider only the entities which are consistent with the overall stance of the article. In particular, journalists often cite the opposing view, and thus can mention a set of characters which is different from the one aligned with stance. For example, in the article in Figure 1, melting ice is a `victim` of rising temperatures, according to the viewpoint of climate activists. However, while the author cites this viewpoint, they do not agree with it, so the corresponding entities are not considered as potential `hero`, `villain`, and `victim`. To summarise, the main characters are the ones framed so by the author/narrator.

(2) We discard characters that either form the backdrop of the story or are used to illustrate a minor (often competing) idea within the main narrative.[11] For example, in the narrative in Figure 1 ANIMALS_NATURE_ENVIRONMENT stakeholders such as melting ice or Artic animals are only used as a battle ground between the climate activists and denialists; the true characters are the parties expressing their opinion regarding them.

(3) The same character can be referred to several times and be represented with several stakeholders. For example, it is common for news stories to mention both climate change regulations (LEGISLATION_POLICIES) and the politicians that propose them (GOVERNMENTS_POLITICIANS. In such cases, instead of adding multiple stakeholders for a character, we choose the one which was more prominent in the context or could be inferred from the other (i.e. LEGISLATION_POLICIES).

(4) We only consider those characters that are active in the narrative plot, rather than references to potential or past heroes, villains, and victims. For example, a news story that paints Republicans as a `villain` for not implementing climate change measures[12] concludes with the following sentence:

> For 2020 and beyond, climate justice will have to become the most animating issue for Democrats.

Since the positive impact of Democrats is only hoped for or predicted to happen in the future, Democrats (or GOVERNMENTS_POLITICIANS) are not an active `hero`, and, overall, the `hero` in this news story is absent. It is important not to assign extraneous entities to the character slots, even if they are otherwise empty, as it will later help to differentiate between narratives. For example, here it allows us to distinguish a narrative criticizing the `villain` from alternative narratives which depict an active conflict between `hero` and `villain`.

(5) For the same reasons, we do not add stakeholders that are only implied but not directly referred to in the text. For example, we do not add ANIMALS_ENVIRORNMENT as a `victim` unless it is specifically mentioned, though it can be inferred from the majority of pro-climate action news stories. Similarly, though the stories warning about the dangers of climate inaction are usually inspired by scientific evidence, SCIENCE_EXPERTS_SCI.REPORTS are not a `hero` in them unless they have an active role, as in here[13]

> Climate report warns of extreme weather, displacement of millions without action

This allows to differentiate between a narrative which appeals to authority of scientists (so called "Gore" narrative) from a similar but often more emotionally charged and less "objective" alarmist narrative ("12 Years to save the world").

C Annotation details

The annotation was performed in two stages:

C.1 Stage 1: annotating Hero, Villain, and Victim

During stage 1, we employed three external annotators, all specializing in social sciences and familiar with Narrative Policy Framework, and one of the authors of the article, who is an expert annotator with knowledge of media discourse and framing. All four annotators were asked to read a set of articles, specify who they think are the main Hero, Villain, and Victim in each of them, and provide some explanations (see an example in Figure 4).

The external annotators were expected to spend around 10 minutes per article, and were compensated at a competitive rate of 35 USD per hour. Each of the 100 articles in the dataset was annotated at least by two external annotators, and all of them were annotated by the internal expert annotator. Since the annotators were asked to specify entities for Hero, Villain, and Victim in free form, their annotations were not directly comparable (e.g. "Biden" vs "Joe Biden" vs "Democrats"). Thus, to evaluate the annotation agreement, as well as to convert the data to a more abstract and useful structure (see Section 3.1), the expert annotator mapped the specific characters mentioned by each of the annotators to their character classes (Appendix A).

We evaluate the agreement between all four annotators using Krippendorf's α, and report the averaged agreement of each of the three external annotators with the expert. For the latter, we use the standard metrics of agreement rate (=accuracy), Cohen's κ, and the less commonly used Gwet's AC1,

[11] We are well aware that news stories are complex in terms of interplay of narratives within them and most of them contain what Fløttum and Gjerstad (2017) refers to as narrative polyphony. We intentionally restrict the task to identification of the main narrative only as the first step in disentangling narrative complexity.

[12] Narrative 512 in our dataset.

[13] Narrative 537 in our dataset.

Which entity did you identify as the HERO of the article (if any)?	Which entity did you identify as the VILLAIN of the article (if any)	Which entity did you identify as the VICTIM of the article (if any)?	Explanation (1-3 sentences)
None	Political Inaction, Italian apathy / lack of concern for climate change	People and City of Venice	public and a politician as essentially blaming a lack of serious climate change action for the floods. It's not clear the prime minister, Conte is being potrayed as HERO, as per the annotated article, maybe it is even portraying him as a villian for not taking sufficient action. Climate change, global warming, and the extreme floods are not
Climate Action, Those taking climate change seriously	Trump	Society/Nature	Article is criticizing politics, specifically the US president Donald Trump for his climate denial or at the very least climate inaction and prioritization of other political agendas. Article quotes leading scientists who warn about the devestating effects of climate change and there call for decisive action.
Other National leaders taking climate change seriously, UN Secretary	Trump	nature/society	Article heavily criticizes Trump for pretending to take climate change/action seriously, while actually not doing very much about it because he is actually a climate sceptic. Article also indicates other world leaders take climate change more seriously and implies Trump should follow suit

Figure 4: Example of stage 1 annotations (Hero, Villain, Victim

	Hero	Villain	Victim
Krippendorf's α	0.757	0.673	0.812
Agreement rate	0.852	0.855	0.927
Cohen's κ	0.783	0.745	0.876
Gwet't AC1	0.837	0.843	0.914

Table 2: Inter-annotator agreement for Hero, Villain, and Victim annotation

	Focus	Conflict	Cultural story
Krippendorf's α	0.780	0.820	0.801
Agreement rate	0.867	0.867	0.867
Cohen's κ	0.776	0.817	0.800
Gwet't AC1	0.810	0.824	0.801

Table 3: Inter-annotator agreement for Focus, Conflict, and Culture story annotation

which compensates for the high imbalance in data distribution. The resulting inter-annotator agreement statistics can be found in Table 2. Overall, we observe acceptable to strong levels of agreement between all four annotators, as well as very high average agreement of each of the annotators with the expert (as judged based on Gwet's AC1). A relatively lower agreement for Hero and Villain in comparison to Victim is explained by the fact that the annotators sometimes chose entities belonging to different stakeholder types to represent the same event. For example, a particular climate initiative can be represented both by a parliamentary bill such as New Green Deal (LEGISLATION_POLICIES), and by the group of people behind it (GOVERNMENTS_POLITICIANS).

C.2 Stage 2: annotating Focus, Conflict, and Culture story

In the second stage, the expert annotator annotated all 100 articles in terms of their Focus, Conflict, and Culture story. Next, a random sample of 30 articles was annotated by another internal annotator who is also an expert in Narrative Policy Framework and framing analysis. The instructions for the annotation and an example of an annotated article are shown on Figures 5 and 6 respectively. To ensure a high quality of the resulting dataset, all disagreements were discussed and adjudicated, and then the corresponding changes were reflected in the samples beyond this calibration study, if necessary.

Table 3 shows the agreement statistics between the two annotators in Stage 2, using the same metrics as for Stage 1. We observe high agreement rates for all three classes, with other scores varying slightly due to number of classes and class distribution, but all being within the strong or very strong agreement range. Disagreement analysis revealed that there were disagreements on Focus between Villain and Victim, when both were discussed at similar length and depth in the article. For Conflict and Cultural story, the disagreement were more systematic (such as confusion between Fuel Resolution and Prevent Conflict, or between Hierarchical and Egalitarian stories); the insights arising from the discussion were reflected in the final labels and allowed us to refine the definitions of these concepts for the prompts used in LLM experiments.

Annotation with vs without narrative frame structure

In this section, we empirically test if structural components help to differentiate between narratives. Specifically, we compare agreement in narrative detection when using a structure-based annotation approach (bottom-up; as described above) vs using a more traditional approach where the annotators are asked to classify narratives top-down based on their descriptions.

For the structure-based approach, we estimate the agreement based on the sample of 30 articles we used for Stage 2 annotation (see Appendix C). In particular, we assume that both annotators agree on a particular narrative if they choose exactly the same values for all its components. For the tradi-

Annotation Instructions

You will be given a full text of an article about climate change and asked to identify some elements of its narrative structure according to Narrative Policy Framework (NFP).

In NFP, a narrative contains at least one of the following characters: **Villain**, who is creating some conflict/problem; **Hero**, who is trying to resolve a conflict or problem; or a **Victim**, who is negatively affected by a conflict or problem. Not all of these characters need to be present in a narrative at the same time.

Each article can contain a mixture of narratives and thus have multiple villain-hero-victim sets. However, we can derive the overarching narrative of the article by determining its main hero, villain, or victim. For each of the articles you will see, we have already annotated the main hero, villain, and victim, so that you can focus on the main narrative characters when you do your annotation.

You will be asked three questions:

1. **Focus**: narratives can have the same characters (hero, villain, victim) but focus on different ones of them. For example, a narrative about negative effects of pollution on environment can focus either on the villain (criticise policies, governments, industries that cause pollution while mentioning its negative effects), or on the victim (describe negative effects on people or nature in detail while also mentioning the culprit).
   Which of the characters (Hero, Villain, or Victim) is the focus of the narrative?

2. **Conflict and Resolution**: apart from their characters, narratives in NFP are defined by the conflict/problem or its resolution described in them. In our case, the conflict/problem is climate change, and resolution is measures against climate change. Thus, a particular narrative can:

   FUEL RESOLUTION: propose or describe specific measures, policies, or events that would contribute to the resolution of the climate crisis.
   FUEL CONFLICT: propose or describe specific measures, policies, or events that would exacerbate the climate crisis.
   PREVENT RESOLUTION: criticise measures, policies, or events that contribute to the resolution of the climate crisis; or deny the climate crisis
   PREVENT CONFLICT: criticise measures, policies, or events that exacerbate the climate crisis; or provides the evidence of climate crisis.

   Please be mindful that the perspective of the author/narrator and the characters in the story regarding the conflict and its resolution can be different; identify and annotate the main perspective which corresponds to the author's/narrator's intention.

   Does this narrative fuel conflict, fuel resolution, prevent conflict, or prevent resolution?

3. **Cultural story**: narratives of climate change are aligned with the following cultural stories, which capture the ideas of the necessity of top-down control vs self-regulation, and the idea of group responsibility vs individual responsibility.

   HIERARCHICAL: this story assumes that the nature can be controlled but we need to be bound by tight social prescriptions. The villain is mismanaged society which led to excessive growth, and heroes are impartial scientists or government intervention.
   INDIVIDUALISTIC: this story assumes that the nature is resilient and will return to equilibrium. Villains here are people who try to control climate change or seek policy changes, and the heroes allow markets to move naturally as individuals compete to create innovative technologies.
   EGALITARIAN: this story assumes that the nature is fragile and there is little opportunity to correct mistakes. The cause of climate change is overconsumption; villains are profit-driven corporations and anyone who supports status quo, and heroes are groups who seek fundamental changes.
   FATALIST: the story assumes that the nature cannot be controlled, and climate change is inevitable whatever efforts we make.

   Which of the cultural stories (Hierarchical, Individualistic, Egalitarian, or Fatalist) does the narrative align with?

Figure 5: Instructions for stage 2 annotation (Focus, Conflict, Culture story

ID: 225

Article:

E.P.A. Plans to Get Thousands of Deaths Off the Books by Changing Its Math Want climate news in your inbox ? Sign up here for Climate Fwd :, our email newsletter. WASHINGTON — The Environmental Protection Agency plans to change the way it calculates the health risks of air pollution, a shift that would make it easier to roll back a key climate change rule because it would result in far fewer predicted deaths from pollution, according to five people with knowledge of the agency 's plans. The E.P.A. had originally forecast that eliminating the Obama - era rule, the Clean Power Plan, and replacing it with a new measure would have resulted in an additional 1,400 premature deaths per year. The new analytical model would significantly reduce that number and would most likely be used by the Trump administration to defend further rollbacks of air pollution rules if it is formally adopted. The proposed shift is the latest example of the Trump administration downgrading the estimates of environmental harm from pollution in regulations. In this case, the proposed methodology would assume there is little or no health benefit to making the air any cleaner than what the law requires. Many experts said that approach was not scientifically sound and that, in the real world, there are no safe levels of the fine particulate pollution associated with the burning of fossil fuels. Fine particulate matter — the tiny, deadly particles that can penetrate deep into the lungs and enter the bloodstream — is linked to heart attacks, strokes and respiratory disease."

Questions:

Considering that in this article the **Villain** is politicians, and the **Victim** is general public, answer the following:

1) Which of the characters (Hero, Villain, or Victim) is the focus of the narrative?

Hero **Villain** Victim

2) Does this narrative fuel conflict, fuel resolution, prevent conflict, or prevent resolution?

Fuels conflict Fuels resolution
Prevents conflict Prevents resolution

3) Which of the cultural stories (Hierarchical, Individualistic, or Egalitarian) does the narrative align with?

Hierarchical Individualistic Egalitarian Fatalist

Figure 6: An example of stage 2 annotation (Focus, Conflict, Culture story

tional approach, we ask two annotators who took part in Stage 1 of annotation (and thus did not classify any elements of the narrative except for its `hero`, `villain`, and `victim`) to choose a narrative for each article based on its description only (as listed in Appendix D).

We find that annotation based on our narrative taxonomy resulted in 63% agreement, while top-down annotation based on the narrative descriptions resulted in a substantially lower 37%. Thus, we can tentatively conclude that structure-based analysis improves narrative detection and understanding. We also observed a reduction in time required for annotation (15 minutes per article based on description of narrative frame vs 7 minutes per article based on its structure, on average).

D Narratives structures and description

In this section we provide a list of discovered narratives, their structures, references to the literature sources where they are mentioned, and definitions taken from that source.

D.1 Narratives focusing on Hero

D.1.1 You're destroying our future

Hero: ENV.ORGS_ACTIVISTS
Villain: GOVERNMENTS_POLITICIANS
Victim: <optional>
Action: FUEL RESOLUTION
Cultural story: EGALITARIAN
Description: The political stasis around climate change means that we cannot rely on politicians to create the change necessary. With collective action, even the politically weak can make a difference and secure a future for generations to come. This can manifest as anything from protests (school strikes) to non-violent civil disobedience.
Source: Bevan (2020)

D.1.2 Technological optimism

Hero: GREEN_TECHNOLOGY_INNOVATION
Villain: INDUSTRY_EMISSIONS, CLIMATE_CHANGE
Victim: <optional>
Action: FUEL RESOLUTION
Cultural story: EGALITARIAN
Description: We should focus our efforts on current and future technologies, which will unlock great possibilities for addressing climate change.
Source: Lamb et al. (2020)

D.1.3 Officials declare climate emergency

Hero: GOVERNMENTS_POLITICIANS
Villain: INDUSTRY_EMISSIONS, CLIMATE_CHANGE, GOVERNMENTS_POLITICIANS
Victim: <optional>
Action: FUEL RESOLUTION
Cultural story: HIERARCHICAL
Description: The climate crisis is sufficiently severe that it warrants declaring a climate emergency. This should occur at different levels of government as climate requires action at all levels, from the hyper-local to the global.
Source: Bevan (2020)

D.1.4 Every little helps

Hero: GENERAL_PUBLIC
Villain: GENERAL_PUBLIC
Victim: <optional>
Action: FUEL RESOLUTION
Cultural story: INDIVIDUALISTIC
Description: This narrative presents a society which has transitioned to a sustainable 'green' way of life. Could be by portraying individuals as the protagonists of stories that propose solutions to climate change.
Source: Bushell et al. (2017)

D.2 Narratives focusing on Villain

D.2.1 12 Years to Save the World

Hero: <optional>
Villain: GOVERNMENTS_POLITICIANS
Victim: ANIMALS_NATURE_ENVIRONMENT, GENERAL_PUBLIC, CLIMATE_CHANGE
Action: PREVENT CONFLICT
Cultural story: HIERARCHICAL
Description: Past and present human action (or inaction) risks a catastrophic future climatic event unless people change their behaviour to mitigate climate change.
Source: Bevan (2020)

D.2.2 Gore

Hero: SCIENCE_EXPERTS_SCI.REPORTS
Villain: GOVERNMENTS_POLITICIANS, GENERAL_PUBLIC, INDUSTRY_EMISSIONS
Victim: ANIMALS_NATURE_ENVIRONMENT, CLIMATE_CHANGE
Action: FUEL RESOLUTION
Cultural story: HIERARCHICAL

Description: This is a narrative of scientific discovery which climaxes on the certainty that climate change is unequivocally caused by humans.
Source: Bushell et al. (2017)

D.2.3 The collapse is imminent

Hero: ENV.ORGS_ACTIVISTS
Villain: GOVERNMENTS_POLITICIANS
Victim: <optional>
Action: FUEL RESOLUTION
Cultural story: EGALITARIAN
Description: The climate crisis is such that some kind of societal collapse is near inevitable. Due to the inaction of the negligent or complacent politicians the social contract has broken down and it is incumbent upon individuals to engage in non-violent civil disobedience to shock society into urgent action.
Source: Bevan (2020)

D.2.4 Climate solutions won't work

Hero: <optional>
Villain: LEGISLATION_POLICIES, GREEN_TECHNOLOGY_INNOVATION
Victim: GENERAL_PUBLIC, ANIMALS_NATURE_ENVIRONMENT
Action: PREVENT RESOLUTION
Cultural story: INDIVIDUALISTIC
Description: Climate policies are harmful and a threat to society and the economy. Climate policies are ineffective and too difficult to implement.
Source: Lamb et al. (2020)

D.2.5 No sticks just carrots

Hero: LEGISLATION_POLICIES
Villain: LEGISLATION_POLICIES
Victim: GENERAL_PUBLIC
Action: PREVENT RESOLUTION
Cultural story: INDIVIDUALISTIC
Description: Society will only respond to supportive and voluntary policies, restrictive measures will fail and should be abandoned.
Source: Lamb et al. (2020)

D.2.6 All talk little action

Hero: <optional>
Villain: GOVERNMENTS_POLITICIANS
Victim: <optional>
Action: PREVENT RESOLUTION
Cultural story: EGALITARIAN
Description: This narrative emphasises inconsistency between ambitious climate action targets and actual actions.
Source: Lamb et al. (2020)

D.2.7 Victim blaming

Hero: <optional>
Villain: GENERAL_PUBLIC
Victim: GENERAL_PUBLIC
Action: PREVENT RESOLUTION
Cultural story: INDIVIDUALISTIC
Description: Individuals and consumers are ultimately responsible for taking actions to address climate change.
Source: Lamb et al. (2020)

D.2.8 Debate and scam

Hero: <optional>
Villain: GOVERNMENTS_POLITICIANS, LEGISLATION_POLICIES, ENV.ORGS_ACTIVISTS, MEDIA_JOURNALISTS
Victim: <optional>
Action: PREVENT RESOLUTION
Cultural story: INDIVIDUALISTIC
Description: The heroes of this narrative are sceptical individuals who dare to challenge the false consensus on climate change which is propagated by those with vested interests.
Source: Lamb et al. (2020)

D.2.9 Others are worse than us

Hero: GOVERNMENTS_POLITICIANS
Villain: GOVERNMENTS_POLITICIANS
Victim: <optional>
Action: PREVENT RESOLUTION
Cultural story: INDIVIDUALISTIC
Description: Other countries, cities or industries are worse than ourselves. There is no point for us to implement climate policies, because we only cause a small fraction of the emissions. As long as others emit even more than us, actions won't be effective.
Source: Lamb et al. (2020)

D.3 Narratives focusing on Victim

D.3.1 Endangered species

Hero: <optional>
Villain: GOVERNMENTS_POLITICIANS, LEGISLATION_POLICIES, INDUSTRY_EMISSIONS
Victim: ANIMALS_NATURE_ENVIRONMENT
Action: PREVENT CONFLICT
Cultural story:HIERARCHICAL
Description: Endangered species (like polar bears) are the helpless victims of this narrative, who are

seeing their habitat destroyed by the actions of villainous humans.
Source: Bushell et al. (2017)

D.3.2 We are all going to die

Hero: <optional>
Villain: CLIMATE_CHANGE, INDUSTRY_EMISSIONS
Victim: GENERAL_PUBLIC
Action: PREVENT CONFLICT
Cultural story: EGALITARIAN
Description: This narrative shows the current or potential catastrophic impact of climate change on people.
Source: Shanahan (2007)

D.3.3 Carbon fueled expansion

Hero: <optional>
Villain: LEGISLATION_POLICIES, GREEN_TECHNOLOGY_INNOVATION
Victim: GENERAL_PUBLIC, INDUSTRY_EMISSIONS
Action: PREVENT RESOLUTION
Cultural story: INDIVIDUALISTIC
Description:The free market is at the centre of this narrative which presents action on climate change as an obstacle to the freedom and well-being of citizens. The narrative can stress social justice or well-being of individual citizens.
Source: Bushell et al. (2017)

E Dataset statistics

In Figure 7 we show the distribution of labels for frame components and the resulting narrative frames.

F Model sizes, costs and parameters

Mixtral-8x7B-Instruct-v0.1 (Mixtral): 46.7B params

gemini-1.5-pro (Gemini): 1.5T params

Llama-3.1-8B-Instruct: 8B params

Approximate experiments costs: 600 USD.

Hyperparameters for Llama LoRA fine-tuning:

max seq length = 4000
r = 16
target modules = "q_proj", "k_proj", "v_proj", "o_proj", "gate_proj", "up_proj", "down_proj"
lora alpha = 16
lora dropout = 0
batch size = 2
gradient accumulation steps = 4
warmup steps = 5
learning rate = 2e-4
optim = adamw8bit
weight decay = 0.01

G Basic prompts

In tables below we show the basic prompts used for the classification: Table 4 for Hero, Villain, Victim and Focus classes, Table 5 for Conflict and resolution classification, Table 6 for Story classes, and Table 7 for Narrative frame classification.

H Modified prompts with structure descriptions

In tables below we show the modified prompts used for Narratives (Table 8).

I Additional experiment details

We examine if the performance can be improved by exposing models to annotated examples and optimizing the prompts by adding Chain-of-Thought steps. First, we use 5 randomly selected samples from our dataset for 5-shot learning with GPT4o model. However, except for Hero stakeholder identification, where it leads to some gains, it causes overgeneralization to seen labels and thus drop in performance (see Table 9). We observe similar effects when we perform Low-Rank Adaption (LoRA) fine-tuning (Hu et al., 2021) of Llama.[14] Similarly, we notice that the fine-tuned model tends to overpredict the most prominent labels, discarding minor classes.

We also use the 5 random samples for a DSPy program (Khattab et al., 2023) to automatically generate and optimize reasoning steps for Chain-of-Thought (CoT) prompting. The gains (compared to non-optimized 5-shot prompting) are also minimal (see Table 9). In addition, we tried implementing Chain-of-Thought (CoT) manually for HVV identification tasks, where we guide the model through the steps of identifying candidate entities, choosing

[14]We choose Llama as a stronger model among open-source ones, and perform 5-fold fine-tuning with 20% holdout set, ensuring balanced class representation (hyperparameters and details in Appendix F): despite improved classification of Hero, the overall performance drops (Table 10).

You are a social scientist specializing in climate change. You will be given a newspaper article and asked who is framed as a hero, villain or a victim in it. For each of these categories, you will be also asked to specify the corresponding word or phrase, and to classify it into the following classes: GOVERNMENTS_POLITICIANS: governments and political organizations; INDUSTRY_EMISSIONS: industries, businesses, and the pollution created by them; LEGISLATION_POLICIES: policies and legislation responses; GENERAL_PUBLIC: general public, individuals, and society, including their wellbeing, status quo and economy; ANIMALS_NATURE_ENVIRONMENT: nature and environment in general or specific species; ENV.ORGS_ACTIVISTS: climate activists and organizations SCIENCE_EXPERTS_SCI.REPORTS: scientists and scientific reports/research CLIMATE_CHANGE: climate change as a process or consequence GREEN_TECHNOLOGY_INNOVATION: innovative and green technologies MEDIA_JOURNALISTS: media and journalists Finally, you need to detect which of the characters (hero, villain, or victim) the news story is focusing on. Please return a json object which consists of the following fields: hero_class: a label for the hero from the list above, or 'None' if the hero cannot be identified. villain_class: a label for the villain from the list above, or 'None' if the villain cannot be identified. victim_class: a label for the victim from the list above, or 'None' if the victim cannot be identified. focus: one of the following - HERO, VILLAIN, VICTIM

Table 4: Basic prompt for Hero, Villain, Victim, and Focus classification

most prominent among them, and finally classifying their stakeholder type, but this lead to worse performance.

Overall, these additional experiments show that the tasks are difficult to meaningfully learn from examples or even through reasoning steps.

J Narrative frame prediction with and without structure

Below we show confusion matrices for GPT4o with a basic prompt vs with a structured prompt and oracle (human-annotated) labels.

K Prompts and results for HVV stakeholder extraction

In this sections we provide prompts we used for multi-step clustering and extraction of stakeholder classes, and well as the list of the resulting classes to be used in HVV classification prompts.

K.1 Prompts

We provide prompts for identifying candidate entities in each speech Table 11, and then clustering them into stakeholder types Table 12.

K.2 Resulting classes

- HEALTHCARE: frontline workers, medical professionals, and institutions directly involved in providing care and combatting the pandemic;
- VULNERABLE_POPULATION: individuals at higher risk of severe illness or death from COVID-19;
- GENERAL_PUBLIC: general public, individuals, communities, and society;
- GOVERNMENT_POLITICIANS: national and regional governments and policymakers;
- BUSINESS_ECONOMY: businesses, workers, and the broader economy;
- SCIENCE_EXPERTS: scientists, researchers, and research institutions;
- FAITH_GROUPS: faith-based organizations;
- PANDEMIC: the virus itself and the pandemic;
- GLOBAL_EFFORTS: international organizations, global collaborations, and efforts to address the pandemic on a worldwide scale.

You are a social scientist specializing in climate change. You will be given a newspaper article and asked to identify how it relates to climate crisis. Assign one of the following classes: FUEL_RESOLUTION: the article proposes or describes specific measures, policies, or events that would contribute to the resolution of the climate crisis. FUEL_CONFLICT: the article proposes or describes specific measures, policies, or events that would exacerbate the climate crisis. PREVENT_RESOLUTION: the article criticises measures, policies, or events that contribute to the resolution of the climate crisis; or it denies the climate crisis. PREVENT_CONFLICT: the article criticises measures, policies, or events that exacerbate the climate crisis; or it provides the evidence for the climate crisis. Please return a json object which consists of the following field: action: one of the following labels: FUEL_RESOLUTION, FUEL_CONFLICT, PREVENT_RESOLUTION, PREVENT_CONFLICT.

Table 5: Basic prompt for Conflict classification

You are a social scientist specializing in climate change. You will be given a newspaper article and asked what is the cultural story reflected in it. You should choose one of the following classes: HIERARCHICAL: this story assumes that the situation can be controlled externally, but we need to be bound by tight social prescriptions and group actions. INDIVIDUALISTIC: this story assumes that the situation cannot be controlled externally, and no group actions are necessary. EGALITARIAN: this story assumes that the situation requires combined efforts and group actions of all members of society. Please return a json object which consists of the following field: story: a label from the classes above.

Table 6: Basic prompt for Cultural story classification

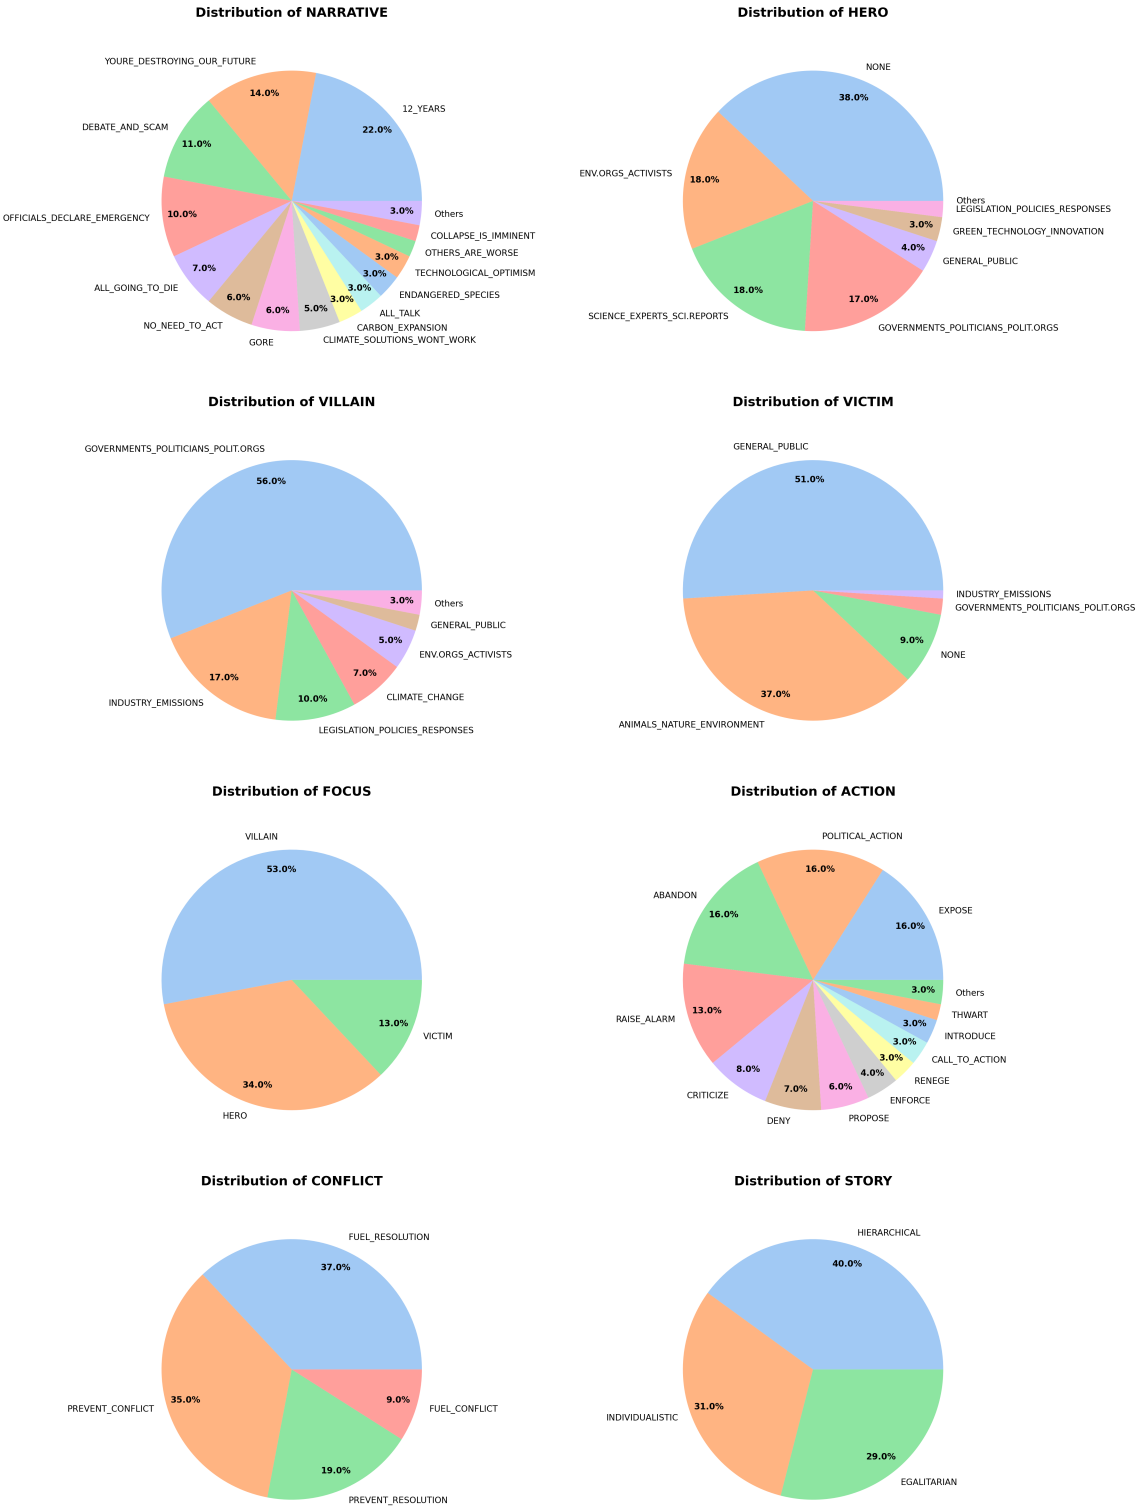

Figure 7: Label distributions for narrative frames and their elements in the dataset

You are a social scientist specializing in climate change.
You will be given a newspaper article and asked what is the main narrative in it.
You should choose one of the following classes:
12_YEARS: 12 Years to save the world - Past and present human action (or inaction) risks a catastrophic future climatic event unless people change their behaviour to mitigate climate change.
ALL_GOING_TO_DIE: We are all going to die - This narrative shows the current or potential catastrophic impact of climate change on people
ALL_TALK: All talk little action - This narrative emphasises inconcistency between ambitious climate action targets and actual actions.
CARBON_EXPANSION: Carbon-fuelled expansion - The free market is at the centre of this narrative which presents action on climate change as an obstacle to the freedom and well-being of citizens.
CLIMATE_SOLUTIONS_WONT_WORK: Climate solutions won't work. Climate policies are harmful and a threat to society and the economy. Climate policies are ineffective and too difficult to implement.
COLLAPSE_IS_IMMINENT: The climate crisis is due to the inaction of the negligent or complacent politicians, and it is incumbent upon individuals to shock society into urgent action
DEBATE_AND_SCAM: The heroes of this narrative are sceptical individuals who dare to challenge the false consensus on climate change which is propagated by those with vested interests.
ENDANGERED_SPECIES: Endangered species (like polar bears) are the helpless victims of this narrative, who are seeing their habitat destroyed by the actions of villainous humans.
EVERY_LITTLE_HELPS: This narrative presents a society which has transitioned to a sustainable 'green' way of life. Could be by portraying individuals as the protagonists of stories that propose solutions to climate change.
GORE: This is a narrative of scientific discovery which climaxes on the certainty that climate change is unequivocally caused by humans.
NO_STICKS: No sticks just carrots - Society will only respond to supportive and voluntary policies, restrictive measures will fail and should be abandoned.
OFFICIALS_DECLARE_EMERGENCY: Officials declare a climate emergency - The climate crisis is sufficiently severe that it warrants declaring a climate emergency. This should occur at different levels of government as climate requires action at all levels, from the hyper-local to the global.
OTHERS_ARE_WORSE: Others are worse than us - Other countries, cities or industries are worse than ourselves. There is no point for us to implement climate policies, because we only cause a small fraction of the emissions. As long as others emit even more than us, actions won't be effective.
TECHNOLOGICAL_OPTIMISM: We should focus our efforts on current and future technologies, which will unlock great possibilities for addressing climate change.
VICTIM_BLAMING: Individuals and consumers are ultimately responsible for taking actions to address climate change.
YOURE_DESTROYING_OUR_FUTURE: The political stasis around climate change means that we cannot rely on politicians to create the change necessary. With collective action, even the politically weak can make a difference and secure a future for generations to come.
Please return a json object which consists of the following field:
narrative: a label from the classes above.

Table 7: Basic prompt for Narrative classification

You are a social scientist specializing in climate change.
You will be given a newspaper article and asked what is the main narrative in it. You should choose one of the following classes:

12_YEARS: 12 Years to save the world - Past and present human action (or inaction) risks a catastrophic future climatic event unless people change their behaviour to mitigate climate change. The villain here is government or industry pollution, and the victim is environment, people, or climate change. The narratives focuses on villain and shows how they deny climate change or abandon climate policies.
ALL_GOING_TO_DIE: We are all going to die - This narrative shows the current or potential catastrophic impact of climate change on people. The villain here is climate change or industry emissions, and the victim is general public. The narrative focuses on victim and raises the alarm.
ALL_TALK: All talk little action - This narrative emphasises inconcistency between ambitious climate action targets and actual actions. The villain here is government and politicians, and the victim is often climate change. The narrative focuses on villain who reneged on their promise to support climate policies.
CARBON_EXPANSION: Carbon-fuelled expansion - The free market is at the centre of this narrative which presents action on climate change as an obstacle to the freedom and well-being of citizens. The villain here is climate policies or green technologies, and the victim is general public or old industries. The narrative focuses on victim and advocates for abandoning climate policies.
CLIMATE_SOLUTIONS_WONT_WORK: Climate solutions won't work. Climate policies are harmful and a threat to society and the economy. Climate policies are ineffective and too difficult to implement. The villain is here climate policies or green technologies, and the victim is usually general public. The narrative focuses on villain and criticizes them.
COLLAPSE_IS_IMMINENT: The climate crisis is due to the inaction of the negligent or complacent politicians, and it is incumbent upon individuals to shock society into urgent action. The heroes here are environmental activists, and the villain is government. The narrative focuses on villain and advocated for taking action such as protests or disobedience.
DEBATE_AND_SCAM: The heroes of this narrative are sceptical individuals who dare to challenge the false consensus on climate change which is propagated by those with vested interests. The villains are governments, activists, journalist and policies that support climate measures. The narrative focuses on villains and exposes them.
ENDANGERED_SPECIES: Endangered species (like polar bears) are the helpless victims of this narrative, who are seeing their habitat destroyed by the actions of villainous humans. The villain here can be government, legislation, industry, and the victim is environment and nature. The narrative focuses on victims and shows how they are endangered.
EVERY_LITTLE_HELPS: This narrative presents a society which has transitioned to a sustainable 'green' way of life. Could be by portraying individuals as the protagonists of stories that propose solutions to climate change. The heroes here are individuals and common people, and it is implied that they are also a villain. The narrative focuses on hero and shows how they change their consumption.
GORE: This is a narrative of scientific discovery which climaxes on the certainty that climate change is unequivocally caused by humans. The heroes here are scientists, the villain is government, general public, or industry pollution, and the victim is environment or climate change. The narrative focuses on villain and raises alarm.
...
Please return a json object which consists of the following field:
narrative: a label from the classes above.

Table 8: Prompt for Narrative classification with Hero, Villain, Victim, and Focus specified (abbreviated)

	Hero	Villain	Victim	Focus	Conflict	Story	Narrative
GPT4o zero-shot	0.325	0.454	0.266	0.656	0.332	0.574	0.258
GPT4o 5-shot	0.414	0.357	0.319	0.613	0.272	0.390	0.190
GPT4o 5-shot with CoT	0.417	0.412	0.330	0.627	0.332	0.430	0.178

Table 9: Macro-averaged F1 performance of GPT4o with 5 shot prompting and Dspy optimization for 7 narrative understanding tasks

	Hero	Villain	Victim	Focus	Action	Story	Narrative
Without LoRA	0.271	**0.156**	**0.336**	**0.568**	**0.379**	**0.449**	**0.181**
With LoRA	**0.338**	0.118	0.221	0.351	0.231	0.393	0.077

Table 10: Macro-averaged F1 performance of Llama 3.1 with vs without LoRA fine-tuning for 7 narrative understanding tasks

You are a social scientist specializing in media analysis. You will be given a politician's address and asked asked who or what is framed as a hero, villain or a victim in it. List the entities corresponding to these character roles, and cluster them according to their type (i.e. what kind of entity they represent). Please return a json object which consists of the following fields: heroes: a list of entity types that you identified as heroes, villains: a list of entity types that you identified as villains, victims: a list of entity types that you identified as victims. Do not include anything apart from these fields.

Table 11: Basic prompt for candidate characters extraction

You are a social scientist specializing in media analysis. You will be given a list of entities that appear in politicians speeches regarding Covid 19. Many of these entities are similar or overlapping. Cluster them to derive the main actors or stakeholders groups.

Table 12: Basic prompt for grouping entities into stakeholder types

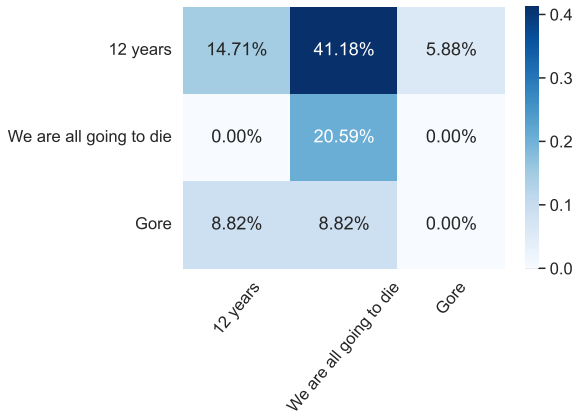

Figure 8: Confusion matrix for zeroshot prediction of only 3 narratives with GPT-4.

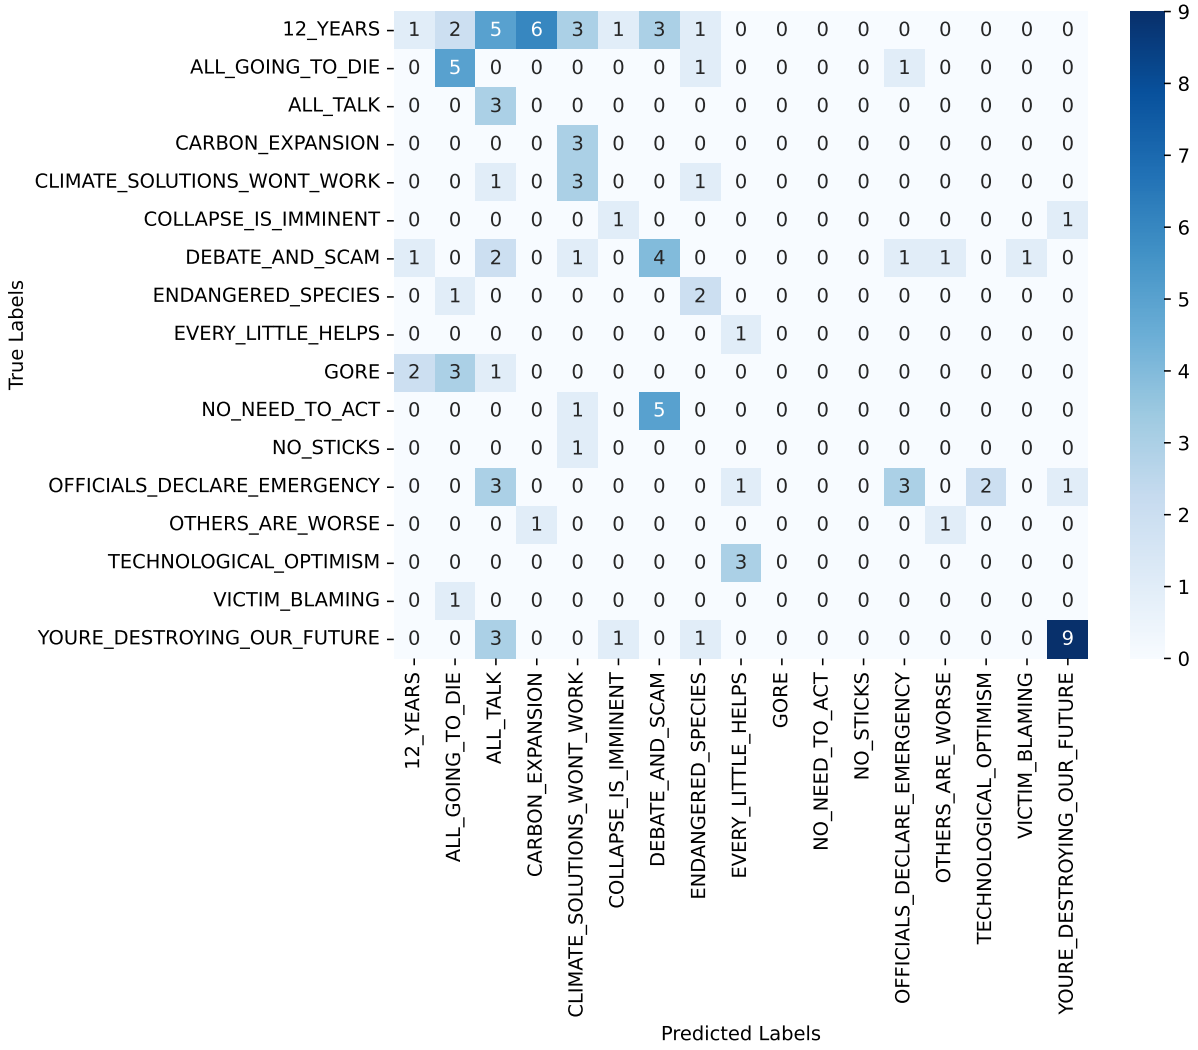

Figure 9: Confusion matrix for Narrative frames prediction using the basic prompt

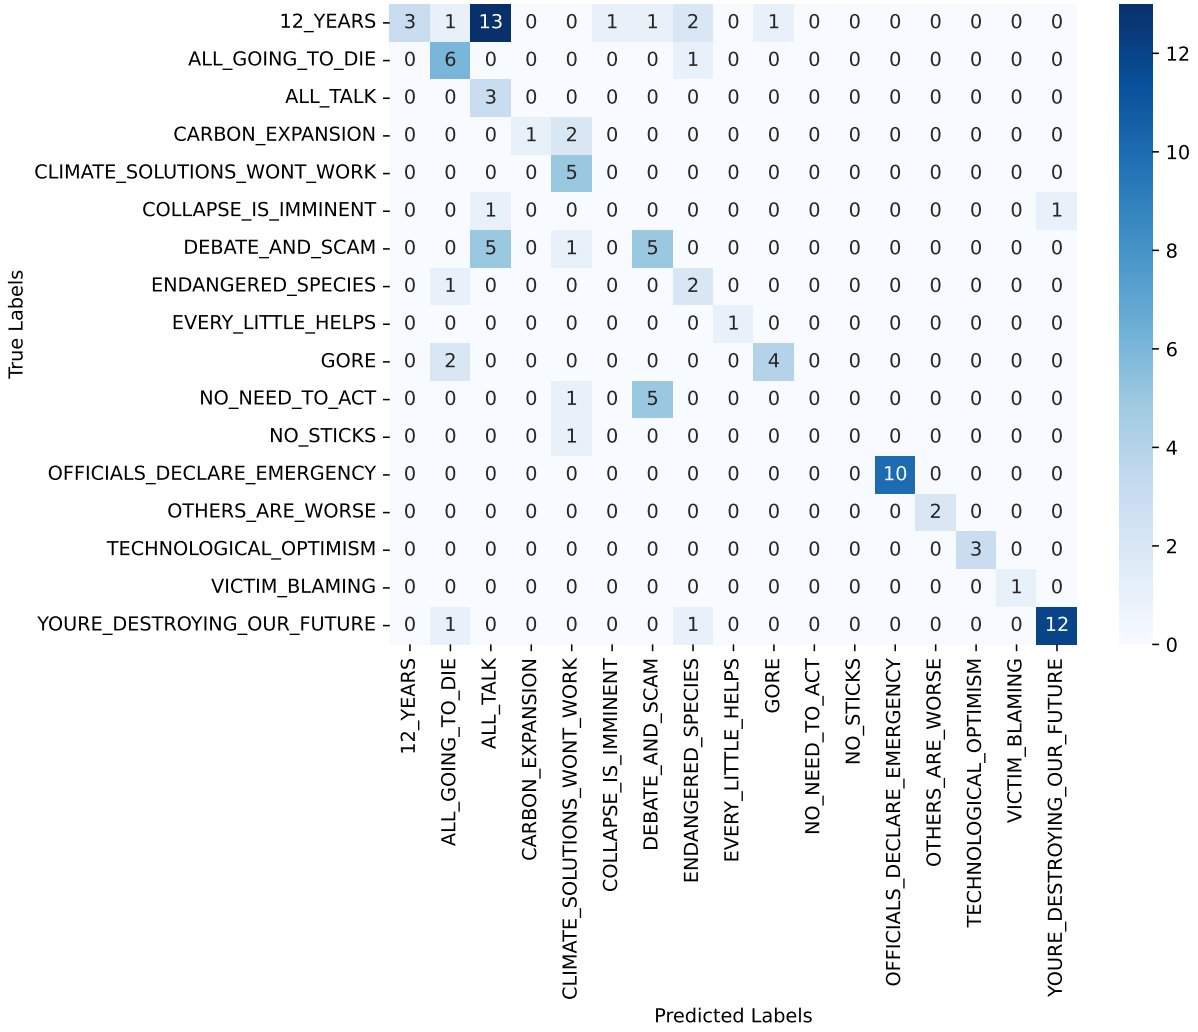

Figure 10: Confusion matrix for Narrative frames prediction using the structured prompt with oracle labels