

# Investigating Controversy Framing across Topics on Social Media

Maxwell A. Weinzierl and Sanda M. Harabagiu

Human Language Technology Research Institute, The University of Texas at Dallas  
{maxwell.weinzierl, sanda}@utdallas.edu

## Abstract

Controversial discourse is abundant on social media. Understanding how controversial problems are framed in online discourse is crucial for gaining insights into public opinion formation and for addressing misinformation and polarization. This paper presents a novel method for discovering and articulating framing of controversial problems, enabling the investigation of how controversy is framed across several diverse topics. The promising results, made possible by recent advances in Large Language Models, indicate that discovering framings across topics is feasible. The discovered frames offer valuable insights into how and why controversial problems are discussed on social media.

## 1 Introduction

Pew Research Center (Pew, 2024) indicates that social media now plays a crucial role in shaping public discourse - especially among younger adults - by serving as a primary source for news and opinion shaping. In this digital public sphere, how issues are understood and opinions formed depends not only on the content that is shared, but also on the way that content is *framed* (Entman, 2003; Reese et al., 2001; Scheufele, 2004; Chong and Druckman, 2012; Bolsen et al., 2014).

As defined by Entman (1993) “to frame is to select some aspects of a perceived reality and make them more salient in a communicating text, in such a way as to promote problem definition, causal interpretation, moral evaluation, and/or treatment recommendation for the item described.” Most previous work (e.g. Card et al. (2016); Naderi and Hirst (2017); Field et al. (2018); Khanehzar et al. (2019); Kwak et al. (2020a); Mendelsohn et al. (2021)) focused **only** on the automatic discovery of framing *aspects* which are *addressed* throughout social media posts. Figure 1 shows how, for the topic of immigration, 27 different aspects/problems (a few illustrated in the Figure) were automatically

detected in posts by Mendelsohn et al. (2021) when the identification of frame aspects *addressed* in posts was cast as a multi-label classification problem.

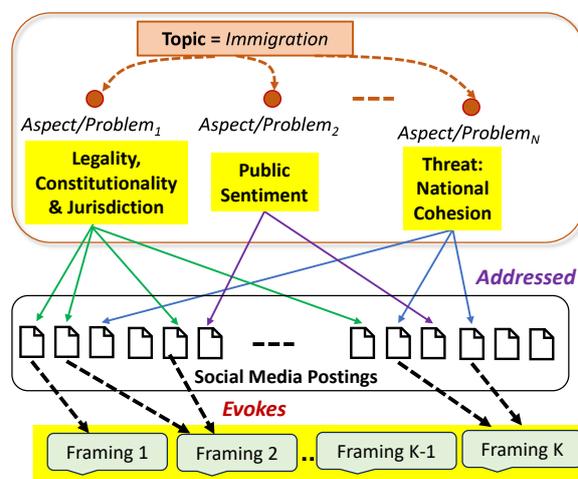


Figure 1: Framing as discovery of Social Media Posts addressing controversial aspects/problems of a topic and generation of Frames of Communication (FoCs) providing interpretations to the controversial topic problems.

Although the identification of the frame aspects addressed in posts is very important, we claim it does not inform us **how** each of those controversial aspects is framed. We know only which reality aspects/problems are highlighted by the frames, but we are not aware of how these controversial aspects are interpreted or framed. Entman’s definition indicates that the causal interpretation of the topic’s problems reveals how the problems are framed. Therefore, we propose to automatically generate the **Frames of Communication (FoCs)** that are (1) *evoked* from posts and (2) provide an interpretation of the framing of the topic’s highlighted aspects/problems. In this way, as shown in Figure 1, following the framework introduced in Weinzierl and Harabagiu (2024a), we consider that (a) each FoC may be **evoked** by multiple posts; and (b) each FoC provides a different *interpretation* of

a topic’s aspect; while (c) multiple posts may **address** the same topic’s aspect. The fact that each topic aspect affords multiple interpretations, provided by the FoCs evoked in the posts that address it, indicates that these topic aspects correspond to *Controversial Problems* (CPs) associated with a topic.

An example of a post that evokes an FoC is illustrated in Figure 2. The same post addresses a CP, for which the FoC provides an interpretation.

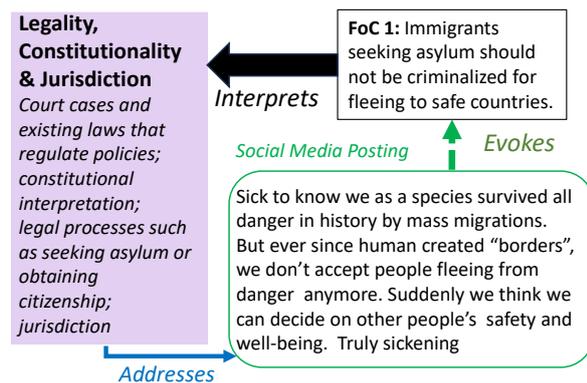


Figure 2: Example of a Frame of Communication evoked in a Social Media Post.

Evidently, discovering FoCs is far more difficult than identifying CPs in posts. First of all, unlike controversial problems, which are known (and their definitions are available), the FoCs are not known, therefore they cannot be extracted from the posts. Instead, the **discovery** of FoCs has to (1) recognize if a post may evoke an FoC; and (2) if so, each FoC must be revealed through its **articulation**; and (3) the CP(s) interpreted by the FoC must be identified.

Given that understanding how controversies are framed in online discourse is crucial for gaining insights into public opinions and for addressing misinformation and polarization (Weinzierl and Harabagiu, 2021, 2022a), we believe it is important to develop automatic methods for the discovery and articulation of FoCs. Awareness of the FoCs would inform the design of more effective counter-narratives against hate speech or inform public interventions against misinformation (Iyengar, 1991; Entman, 2003; Weinzierl and Harabagiu, 2022b).

A pioneering method for the discovery and articulation of FoCs was reported in Weinzierl and Harabagiu (2024a). This method was successful in discovering and articulating FoCs and the CPs they interpret by combining Chain-of-Thought prompting of Large Language Models (LLMs) with In-Context Active Curriculum Learning. Notably,

the method reported in Weinzierl and Harabagiu (2024a) operated by *processing each post separately*, consolidating the discovered FoCs at a later stage, therefore, operating within the framework illustrated in Figure 1 in a **Bottom-up manner**. However, this method has two limitations: (1) it was tested only on one topic: Hesitancy for COVID-19 vaccination, and (2) it required significant human effort due to the active learning framework on which it relied. Therefore, it is an open question whether a method can be developed that can both generalize *across* topics and require no human-in-the-loop to discover and articulate FoCs.

To answer this question, in this paper we introduce a new method that discovers and articulates FoCs evoked in a corpus discussing a topic  $T_2$ , namely  $\mathcal{N}(T_2)$ . Unlike the method reported in Weinzierl and Harabagiu (2024a), our method operates in a **top-down manner**, by *processing the entire corpus*  $\mathcal{N}(T_2)$  *at once* and taking advantage of:

- 1□ an existing reference corpus,  $\mathcal{R}(T_1)$ , discussing  $T_1$ , a different topic. The corpus  $\mathcal{R}(T_1)$  contains annotations of (a) the FoCs evoked in each post and of (b) the CP(s) interpreted by each FoC.
- 2□ the recent capability of LLMs of processing very large contexts; and
- 3□ a novel Corpus-Wide One-Shot (CWOS) prompting method that considers the entire corpus  $\mathcal{R}(T_1)$  and its annotations as a single demonstration, enabling one-shot in-context learning.

We note that this novel top-down method for discovering and articulating FoCs allows us to investigate how controversy is framed when discussing a new topic, given awareness of the framing of controversy in a known topic. Consequently, controversy analysis across topics is made possible. To invite the NLP research community to further explore controversy framing, we make all code, annotations, articulated frames, and interpreted controversial problems available on GitHub <sup>1</sup>.

## 2 The Datasets

We relied on five different datasets: two reference datasets and three additional datasets used for testing and evaluation of frame discovery and articulation method introduced in this paper. The first reference dataset covers the topic of COVID-19 vaccines (hereafter, **C19**), and the second one covers immigration (hereafter, **IMM**). Both datasets

<sup>1</sup><https://github.com/Supermaxman/cross-frame>

include social media posts from Twitter / X along with the annotated frames of communication they evoke. Additionally, for each frame, annotations are available of the underlying problems interpreted and definitions for those problems. The three testing and evaluation datasets consist of posts from Twitter / X discussing the controversial topics of Abortion (AB), Climate Change (CC), and Feminism (FM).

**Dataset for COVID-19 Vaccines (C19):** consists of the only available dataset of social media posts annotated with articulated frames, namely COVAXFRAMES (Weinzierl and Harabagiu, 2022b). This dataset includes frames related to COVID-19 vaccination hesitancy. Vaccine hesitancy is characterized by seven factors, or problems, that increase or decrease an individual’s likelihood of getting vaccinated (Geiger et al., 2021). These problems, along with their definitions, are provided in Appendix A. For the present study, we use the test portion of COVAXFRAMES - which includes 2,815 posts from Twitter / X with 113 evoked frames and the corresponding problem interpretations.

**Dataset for Immigration (IMM):** extends the dataset introduced by Mendelsohn et al. (2021), which consists of 2,730 posts from Twitter / X annotated with any of the 27 generic, narrative, and immigration-specific problems they address. The analysis of immigration framing in Mendelsohn et al. (2021) relied on the 14 issue-generic frame problems introduced by the Policy Frames Codebook (Boydston et al., 2018) and the Media Frames Corpus (Card et al., 2015a), and expanded these problems by introducing 11 issue-specific problems (Benson, 2013a; Hovden and Mjelde, 2019a) and 2 narrative-focused problems (Iyengar, 1990) surrounding the topic of immigration. The definition for each problem is provided in Appendix A. We have further annotated the 72 FoCs evoked in the posts of this dataset. The FoCs were produced by communication expert-led frame analysis, in the same way as in Weinzierl and Harabagiu (2022b).

**Abortion (AB), Climate Change (CC), and Feminism (FM):** The SemEval-2016 Task 6 dataset (Mohammad et al., 2016) was widely employed to train and evaluate automatic stance detection methods on five controversial topics (Sobhani et al., 2016; Du et al., 2017; Li and Caragea, 2019; Xu et al., 2020; Allaway et al., 2021). Each post, sourced from Twitter / X, is annotated with the stance expressed towards one of five controversial

topics: abortion, atheism, climate change, feminism, and Hillary Clinton. The topics of abortion, climate change, and feminism were selected for our study due to extensive prior work investigating the controversial problems addressed for these topics. Appendix A lists these controversial problems for each topic, along with their definitions and sources of the definitions. We utilize all social media posts from each topic, as provided by TweetEval (Barbieri et al., 2020), which consists of 933 posts about Abortion (AB), 564 posts discussing Climate Change (CC), and 949 posts about Feminism (FM). It is important to note that these three test datasets, AB, CC, and FM do not have available annotations of FoCs evoked in their posts, and thus can only serve as evaluation datasets.

### 3 The Method

Our top-down approach for discovering Frames of Communication (FoCs) was inspired by how communication experts perform framing analysis (Reese, 2007; Matthes and Kohring, 2008; Van Gorp, 2010; Russell Neuman et al., 2014). Framing analysis consists of a complex series of inductive, then deductive judgments on an entire corpus, from which expert inference of the articulation of FoCs emerges (Van Gorp, 2010; Walter and Ophir, 2019; Vreese, 2005).

In the inductive phase, the discovery and articulation of FoCs results from expert reasoning with the CPs addressed across multiple texts derived from corpus-wide inspection to arrive at a consistent, causal interpretation of each CP. The deductive phase identifies the texts from the entire corpus where the FoCs are *evoked*, revealing the *relevance* of each FoC, cf. (Gamson, 1989). FoC discovery and articulation is very difficult, even for trained communication experts who rely on codebooks emerging from their reasoning and painful inspection of large quantities of texts (Kwak et al., 2020b; Russell Neuman et al., 2014; Reese, 2007; Matthes and Kohring, 2008).

To our knowledge, corpus-wide automatic discovery and articulation of FoCs (corresponding to the inductive phase) has not been attempted before. However, methods for automatic identification of FoC evocation (corresponding to the deductive phase) have been reported in (Weinzierl and Harabagiu, 2021, 2022a). Therefore, we decided that it is crucial for our new top-down approach to focus *only* on the automatic discovery and articu-

lation of FoCs. For this purpose, we designed the **Corpus-Wide Discovery & Articulation of FoCs (CWDA-FoC)** method, which leverages some recent advances made possible by newer LLMs, answering the following questions.

*Q1: Why did recent LLMs facilitate Corpus-Wide Discovery and Articulation of FoCs?* First, *long-context LLMs*, such as GPT-4o (Wu et al., 2024), o1 (OpenAI et al., 2024), and o3-mini (OpenAI, 2025) allow us to encode entire corpora of social media posts in their contexts, unlike LLMs with smaller context sizes. For the datasets introduced in Section 2, we require LLMs capable of encoding contexts that approach 100,000 tokens - significantly more than models such as BERT (512) (Devlin et al., 2019), LLaMA 2 (4,096) (Touvron et al., 2023), and even GPT-3.5 (16,385) (Ouyang et al., 2022) and earlier versions of GPT-4 (32,768) (OpenAI, 2023).

Second, newer LLMs allow *structured prompting*. Strict structured prompting, which arose as a technique that enables a strict format to be specified for LLM outputs (Zheng et al., 2023; Yang et al., 2023), is ideal for specifying (1) how FoCs are expected to be articulated and (2) what format demonstrations presented to the LLM should follow. The strict output format is imposed using a constrained decoding methodology, which ensures the LLM generates text in a specified format, such as JSON. Constrained decoding deterministically changes the next-token probabilities of an LLM, such that all non-acceptable tokens (based on the specified format) are assigned probability zero.

Third, LLMs must also support long *output* contexts that are sufficient to articulate potentially hundreds of discovered frames, while also supporting strict structured prompting.

*Q2: How can we take advantage of existing corpora annotated with CPs and FoCs for discovering and articulating FoCs from new corpora?* In Section 2 we introduced two corpora of social media posts annotated with the addressed CPs and the evoked FoCs, namely the **C19** corpus and a new one, introduced by this paper, namely **IMM**. Any of these corpora can be considered as possible examples of a Reference Corpus. Formally, a reference corpus  $\mathcal{R}(T_1) = \{p_1^1, p_2^1, \dots, p_m^1\}$  of posts, discussing a topic  $T_1$ , has annotations from (1) a list of CPs:  $\{c_1^1, c_2^1, \dots, c_q^1\}$  and their definitions; (2) a list of articulated FoCs  $\{f_1^1, f_2^1, \dots, f_x^1\}$ , with (3) additional information linking each FoC  $f_i^1$  to the

CP (or CPs)  $c_j^{i,1}$  (or  $\{c_j^{i,1}, \dots, c_k^{i,1}\}$ ) it is interpreting. This formalization of the Reference Corpus informs the design of novel **Corpus-Wide One-Shot (CWOS) prompting** of LLMs to enable our CWDA-FoC method.

The CWOS prompting method, illustrated in Figure 3, uses a prompt that consists of three different contexts:

- 1□ a *Reference Context*  $R_C$ ;
- 2□ a *Discovery Context*  $D_C$ ; and
- 3□ a *Test Context*  $T_C$ .

To generate  $R_C$  a *Context Template* is used, having slots for: (a) the *topic name*; (b) the *name of each CP* and their *definitions*; as well as (c) the corpus of posts. The Context Template is filled with information from the reference collection of posts,  $\mathcal{R}(T_1)$  by a function  $\mathcal{F}_R$ . The function  $\mathcal{F}_R$  fills the slot for the *topic name* with the name of  $T_1$ . Similarly,  $\mathcal{F}_R$  fills the slots for the name and definition of each CP of  $\mathcal{R}(T_1)$  with the name of each  $c_i^1$  and its definition. The slot of the corpus of posts is filled with the text from  $\{p_1^1, p_2^1, \dots, p_m^1\}$ .

The reference corpus  $\mathcal{R}(T_1)$  also informs  $D_C$ , through a *Discovery Template*, having slots of (a) each FoC discovered and articulated from  $\{p_1^1, p_2^1, \dots, p_m^1\}$ , the posts of  $\mathcal{R}(T_1)$ ; as well as (b) sub-slots for CPs interpreted by each FoC. The Discovery Template is filled with information from  $\mathcal{R}(T_1)$  by a function  $\mathcal{F}_D$ . The function  $\mathcal{F}_D$  fills each FoC slot with a  $f_i$  from  $\mathcal{R}(T_1)$ . It also fills the sub-slots for the CPs interpreted by each FoC  $f_i$  with  $\{c_j^i, \dots, c_k^i\}$  if  $f_i$  interprets all those CPs. It is important to note that  $D_C$  does not contain any information indicating in which posts any FoC  $f_i$  from  $\mathcal{R}(T_1)$  is evoked.

When a new corpus of posts discussing another topic  $T_2$ ,  $\mathcal{N}(T_2)$  is considered, we formalize it as  $\mathcal{N}(T_2) = \{p_1^2, p_2^2, \dots, p_n^2\}$ , where each  $p_j^2$  is a post from the  $\mathcal{N}(T_2)$  corpus. However, this test corpus has no annotations of FoCs. But, we are aware that when discussing the topic  $T_2$ , the list of CPs that are addressed are:  $\{c_1^2, c_2^2, \dots, c_n^2\}$ . To discover and articulate FoCs from  $\mathcal{N}(T_2)$ ,  $T_C$  is generated by using again the *Context Template*. This time the Context Template is populated by the function  $\mathcal{F}_R$  operating on  $\mathcal{N}(T_2)$ , thus (a) filling the topic name slot with the name of  $T_2$ ; (b) filling the slots for the name and definition of each CP of  $\mathcal{N}(T_2)$  with their name of each  $c_i^2$  and its definition; and (c) the slot of the corpus of posts is filled with  $\{p_1^2, p_2^2, \dots, p_n^2\}$ .

As shown in Figure 3, the prompt consisting of a concatenation of  $R_C$ ,  $D_C$ , and  $T_C$  is presented

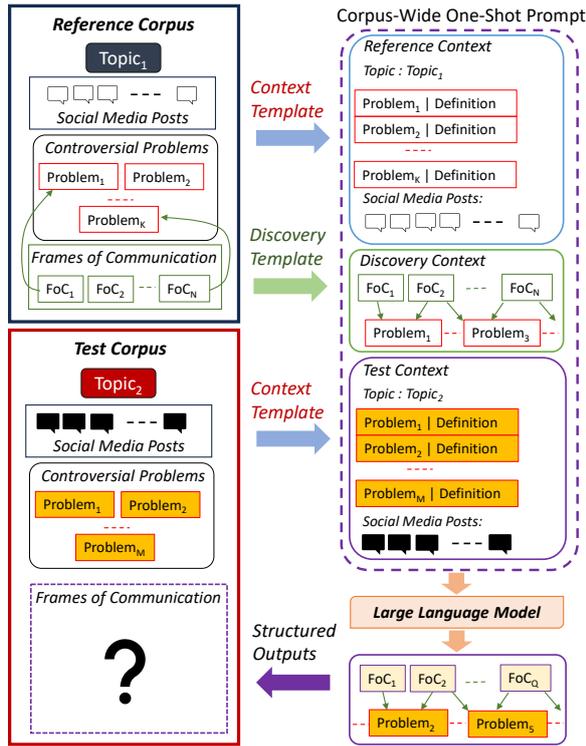


Figure 3: Corpus-Wide One-Shot (CWOS) prompting for FoC discovery and articulation.

to an LLM. While the  $R_C$  provides the necessary information to perform FoC discovery and articulation,  $D_C$  is key in teaching the LLM how it should perform this complex task, and together  $R_C$  and  $D_C$  serve as a single one-shot demonstration.  $T_C$  then invites the LLM to (1) generate  $\{f_1^2, f_2^2, \dots, f_y^2\}$ , the FoCs evoked in  $\mathcal{N}(T_2)$ ; and (2) to provide information linking each FoC  $f_i^2$  to the CP (or CPs)  $c_j^{i,2}$  (or  $\{c_j^{i,2}, \dots, c_r^{i,2}\}$ ) it is interpreting.

In summary, the CWDA-FoC method relies on prompting new LLMs using the CWOS prompting method when (a) considering a Reference Corpus discussing a topic  $T_1$ ; and (b) discovering the FoCs from a New Corpus discussing a different topic  $T_2$ . Figure 4 illustrates the CWDA-FoC method and the details of the CWOS prompting method when using as Reference Corpus **C19** and the new corpus **IMM**. All details of the CWOS prompting methods are presented in Appendix B.

#### 4 Evaluation Results

We considered the only two methods proposed for FoC discovery and articulation: a Bottom-Up approach, and our top-down Corpus-Wide Discovery and Articulation of Frames of Communication (CWDA-FoC) approach, introduced in Section 3.

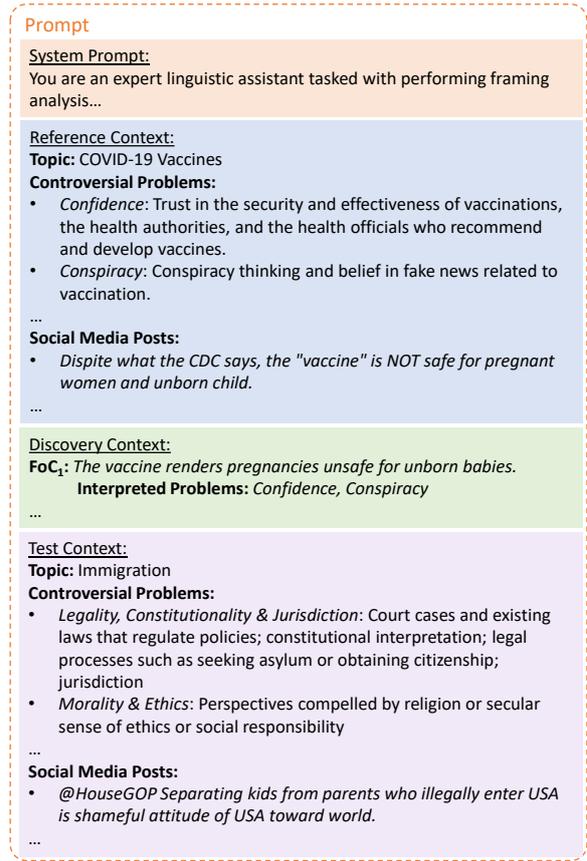


Figure 4: Example of the Corpus-Wide One-Shot (CWOS) prompting framework for frame discovery and articulation.

Prior work introduced CoT-ICACL (Weinzierl and Harabagiu, 2024a), a Bottom-Up FoC discovery and articulation method that operates solely on a single topic. CoT-ICACL leverages a human-in-the-loop active learning paradigm along with in-context curriculum learning to have an expert produce human-verified CoT rationales and demonstrations. However, both the reference demonstrations and test posts must discuss the same controversial topic, restricting this approach significantly.

In contrast, our CWDA-FoC method uses a single one-shot demonstration derived from an external reference corpus of posts and the FoCs they evoke, applied to an entirely separate test corpus with a different topic. CWDA-FoC also requires

Method	Model	Final Frames	$Z$	$A$	$R$	$R_K$	$F_1$	$P_A$
<b>Reference Topic: COVID-19 Vaccines</b>								
CoT-ICACL	LLaMA-2-70B	340	35.29	68.86	42.06	47.32	52.22	42.11
CoT-ICACL	GPT-3.5	386	39.38	53.37	89.57	78.76	66.88	39.39
CoT-ICACL	GPT-4	292	97.60	95.89	94.92	86.73	95.40	93.81
<b>Reference Topic: Immigration</b>								
CWDA-FoC	Gemini 2.0 Flash	67	70.15	64.18	30.28	12.39	42.30	54.72
CWDA-FoC	o1	70	84.29	82.86	38.67	18.58	53.01	75.51
CWDA-FoC	Gemini 1.5 Pro	68	94.12	92.65	41.18	20.35	57.29	88.89
CWDA-FoC	GPT-4o	117	72.65	65.81	51.33	35.40	60.16	48.05
CWDA-FoC	Gemini 2.0 Pro	79	88.61	91.14	46.15	25.66	60.69	86.00
CWDA-FoC	o3-mini	83	98.80	98.80	46.86	17.70	63.57	98.41

Table 1: Evaluation results for bottom-up and top-down methods for discovering and articulating Frames of Communication from social media posts discussing COVID-19 Vaccines from **C19**.

Model	Final Frames	$Z$	$A$	$R$	$R_K$	$F_1$	$P_A$
<b>Reference Topic: COVID-19 Vaccines</b>							
Gemini 2.0 Flash	115	41.74	46.09	47.32	18.06	44.36	39.22
GPT-4o	166	37.35	44.58	58.73	27.78	45.66	36.99
o1	66	59.09	66.67	41.90	15.28	49.04	60.00
Gemini 1.5 Pro	134	62.69	67.91	63.64	27.78	63.16	62.28
Gemini 2.0 Pro	97	72.16	76.29	61.67	36.11	66.50	67.61
o3-mini	249	95.18	99.60	97.25	90.28	96.21	99.46

Table 2: Evaluation results for the Corpus-Wide Discovery and Articulation of Frames of Communication (CWDA-FoC) from social media posts discussing immigration from **IMM**.

Model	Final Frames	$Z$	$A$
<b>Reference Topic: COVID-19 Vaccines</b>			
Gemini 2.0 Flash	179	12.29	9.50
GPT-4o	106	70.75	72.64
Gemini 1.5 Pro	43	79.07	74.42
Gemini 2.0 Pro	87	85.06	80.46
o1	90	86.67	85.56
o3-mini	351	92.59	92.31
<b>Reference Topic: Immigration</b>			
Gemini 2.0 Flash	122	32.79	22.95
GPT-4o	52	75.00	42.31
Gemini 2.0 Pro	46	73.91	69.57
o1	117	74.36	80.34
Gemini 1.5 Pro	37	89.19	91.89
o3-mini	75	96.00	94.67

Model	Final Frames	$Z$	$A$
<b>Reference Topic: COVID-19 Vaccines</b>			
Gemini 2.0 Flash	100	9.00	8.00
Gemini 2.0 Pro	76	23.68	23.68
Gemini 1.5 Pro	24	37.50	41.67
GPT-4o	37	75.68	62.16
o1	48	68.75	70.83
o3-mini	50	88.00	84.00
<b>Reference Topic: Immigration</b>			
Gemini 2.0 Flash	209	2.87	2.39
Gemini 2.0 Pro	25	24.00	20.00
Gemini 1.5 Pro	31	35.48	32.26
GPT-4o	50	36.00	34.00
o3-mini	82	51.22	47.56
o1	68	76.47	75.00

Table 3: Evaluation results for the Corpus-Wide Discovery and Articulation of Frames of Communication (CWDA-FoC) from social media posts discussing abortion from **AB**.

Table 4: Evaluation results for the Corpus-Wide Discovery and Articulation of Frames of Communication (CWDA-FoC) from social media posts discussing climate change from **CC**.

an LLM to support (1) strict structured outputs, (2) massive context sizes (up to 100,000 tokens), and (3) long output context support, to generate significant numbers of FoCs from thousands of posts. Therefore, we considered various LLMs which satisfied these constraints, including Gemini 1.5 Pro (Team et al., 2024a), Gemini 2.0 Flash and Pro

(Team et al., 2024b), GPT-4o (Wu et al., 2024), o1 (OpenAI et al., 2024), and o3-mini (OpenAI, 2025).

For each method, two expert linguists evaluated the final FoCs produced from the social media test datasets, introduced in Section 2. These experts were asked to perform the same judgment process

Model	Final Frames	$Z$	$A$
<b>Reference Topic: COVID-19 Vaccines</b>			
Gemini 2.0 Flash	80	11.25	7.50
Gemini 2.0 Pro	76	35.53	36.84
Gemini 1.5 Pro	45	44.44	42.22
GPT-4o	113	65.49	64.60
o1	207	88.89	87.92
o3-mini	89	92.13	93.26
<b>Reference Topic: Immigration</b>			
Gemini 2.0 Flash	95	9.47	8.42
Gemini 2.0 Pro	86	11.63	10.47
Gemini 1.5 Pro	33	63.64	57.58
GPT-4o	51	78.43	74.51
o3-mini	73	87.67	86.30
o1	80	97.50	98.75

Table 5: Evaluation results for the Corpus-Wide Discovery and Articulation of Frames of Communication (CWDA-FoC) from social media posts discussing feminism from **FM**.

as introduced by Weinzierl and Harabagiu (2024a), such that our new method could be compared to prior work. The experts judged each FoC on two primary dimensions. First, they assessed the **soundness** of the FoCs by asking: Do the FoCs address each of the interpreted problems, and are these interpreted problems correct? Second, experts evaluated the **clarity** of the FoCs by asking: Is the FoC clear and easy to understand, and does it articulate a causal interpretation of the problems?

Using these judgments, we computed the following metrics, which were introduced in Weinzierl and Harabagiu (2024a). The quality of reasoning, denoted by  $Z$ , is defined as the proportion of final FoCs judged to be sound ( $N_S$ ) out of the total number of FoCs produced ( $N_T$ ), i.e.,  $Z = N_S/N_T$ . Similarly, the clarity of the FoC articulation, denoted by  $A$ , is computed as  $A = N_C/N_T$ , where  $N_C$  is the number of FoCs judged clear. Tables 1, 2, 3, 4, and 5 report these results across topics.

When a reference set of FoCs is available, the experts performed an additional judgment: they determined whether a discovered FoC paraphrases any of the provided reference FoCs (i.e., whether it interprets the same problems with the same causal interpretation as a reference FoC). If so, the FoC is marked as *known* (with count  $N_K$ ) out of a total of  $N_F$  reference FoCs. This additional judgment enabled us to compute further metrics: the recall of clearly articulated FoCs,  $R = N_C/(N_C + N_F - N_K)$ ; the known recall,  $R_K = N_K/N_F$ ; a combined  $F_1$  measure defined as  $F_1 = 2AR/(A + R)$ ; and the clarity of novel

FoCs,  $P_A = (N_C - N_K)/(N_T - N_K)$ . Tables 1 and 2 illustrate these judgments for the topics of COVID-19 vaccines and immigration, respectively. The agreement of judgments between linguists was measured with a Cohen’s Kappa score of 0.76, indicating strong agreement (McHugh, 2012).

## 5 Discussion

**Discussion of the Quantitative Evaluation:** The advantages of CWDA-FoC are apparent as (1) it is the first and only method capable of performing cross-topic FoC discovery and articulation, and (2) CWDA-FoC performs extremely well across multiple topics, as illustrated in Tables 1, 2, 3, 4, and 5. On the COVID-19 vaccines dataset, the CWDA-FoC method, as well as the prior CoT-ICACL approach, achieve strong scores on soundness and clarity, as provided in Table 1. However, we see that  $P_A$ , the measure of quality on *new*, never-before found FoCs is much higher for CWDA-FoC. More notably, when evaluated on the immigration dataset, CWDA-FoC achieves high recall and known recall metrics (see Table 2), particularly when paired with models such as o3-mini. The difference between o3-mini’s performance on **C19** vs **IMM** can be explained by the quality and size of the reference data - and therefore the quality of the single one-shot demonstration. The **C19** dataset has 113 evoked FoCs across 2,815 posts, while the **IMM** dataset only contains 72 FoCs evoked across 2,730 posts. The CWDA-FoC results indicate that when the one-shot demonstration used for prompting is sufficiently large and high-quality, it discovers high-quality FoCs evoked for entirely new topics.

The CoT-ICACL method achieved high performance by leveraging multiple chain-of-thought demonstrations for the same COVID-19 Vaccine topic; however, its applicability is limited to scenarios where the reference demonstrations and test posts discuss the same controversial topic. In contrast, the top-down CWDA-FoC method overcomes this limitation by employing a single one-shot demonstration drawn from a reference corpus discussing a *different* topic. CWDA-FoC thereby offers greater flexibility and scalability across topics. We also found that, interestingly, some of the best-performing systems were LLMs trained to perform *reasoning*, such as o1 and o3-mini. We attribute this improved performance to the one-shot nature of our task, relying on a single demonstra-

tion. This indicates that reasoning models have more of an opportunity to “think” about this single demonstration.

For topics where a comprehensive reference set of FoCs is unavailable (abortion, climate change, and feminism as shown in Tables 3, 4, and 5), the evaluation focused solely on soundness and clarity, showing that that CWDA-FoC is effective in articulating evoked FoCs. Although performance varies with the underlying LLM, models like o3-mini and o1 consistently produce FoCs that are both sound and clear, underscoring the robustness of CWDA-FoC across a multitude of controversial topics. We also find that the reference topic of immigration produced more sound and clearly articulated FoCs on the similarly political topics of abortion and feminism, as shown in Table 3 and Table 5 respectively. Alternatively, the reference topic of COVID-19 vaccines produced the best performance on the similarly controversial scientific topic of climate change, as seen in Table 4.

**Discussion of the Qualitative Evaluation:** The analysis of the FoCs articulated by the best-performing CWDA-FoC configurations on each evaluation dataset (AB, FM, and CC) reveals insights into the quality of FoCs automatically discovered on these topics. When immigration was used as the reference topic, o3-mini articulated for the topic of abortion  $FoC_1$ : “*Fetuses do not consent to death, and therefore abortion is murder.*”  $FoC_1$  was also identified by o3-mini as interpreting the problems of “*Fetal Rights*”, “*Sanctity of Life*”, and “*Language & Terminology*”. For the topic of feminism, o1 articulated  $FoC_2$ : “*Accusations of ‘Feminazis’ or ‘ugly feminists’ in media are used to dismiss women’s issues as trivial or extremist.*”  $FoC_2$  was noted as interpreting the problem of “*Media Representation and Backlash*” by o1. When the topic of COVID-19 vaccines was used as a reference topic, o3-mini articulated for the topic of climate change  $FoC_3$ : “*Technical discussions on biodiversity and habitat loss, though data-driven, can render climate change an abstract and distant concern for the average person*” and  $FoC_4$ : “*Unseasonably cool July temperatures are invoked to question long-term climate trends, feeding claims that climate science remains uncertain and abstract.*” The o3-mini model also identified that both  $FoC_3$  and  $FoC_4$  interpret the problem of “*Psychological Distance & Abstraction*”, while  $FoC_4$  also interprets the problem of “*Scientific Uncertainty & Climate Denial*”.

## 6 Related Work

Early research on framing in social media relied heavily on unsupervised approaches (Neuman et al., 2014; Meraz and Papacharissi, 2013; de Saint Laurent et al., 2020) that uncovered lexical patterns and network structures suggestive of framing. While these methods revealed interesting trends, they did not articulate the underlying FoCs or identify the specific problems that these FoCs address. Other studies have developed classifiers to detect frame-invoking language (Baumer et al., 2015), yet these approaches did not capture the causal interpretations embedded within FoCs. A growing body of work has applied supervised NLP methods, often leveraging the Media Frames Corpus (MFC) (Card et al., 2015b), to detect frame-interpreted problems. Techniques ranging from logistic regression (Card et al., 2016) and recurrent neural networks (Naderi and Hirst, 2017) to lexicon induction (Field et al., 2018) and fine-tuning pre-trained language models (Khanehzar et al., 2019; Kwak et al., 2020a) have been used to analyze news and social media. More recently, weakly-supervised methods have been employed to extract subcategories of policy frame dimensions (Roy and Goldwasser, 2020). Notably, Mendelsohn et al. (2021) identified immigration policy problems in social media posts using multi-label classification based on RoBERTa (Liu et al., 2019). However, all of these prior methods focused solely on discovering frame problems rather than articulating them with explicit causal rationales.

More recent research using LLMs has started to expand beyond frame problems. Weinzierl and Harabagiu (2024b) utilized existing FoCs to perform counterfactual reasoning with LLMs to perform zero-shot stance detection, while Weinzierl and Harabagiu (2024c) demonstrated that stance detection towards FoCs has both theoretical and empirical benefits over stance detection towards topics or claims. Islam and Goldwasser (2025) uncovered latent arguments from posts on social media using LLMs, which constructed themes (problems) and talking points for each theme.

## 7 Conclusion

This paper presents a new, top-down method capable of discovering and articulating frames of communication from social media. By leveraging a single one-shot demonstration drawn from an external reference dataset, Corpus-Wide One-Shot (CWOS) prompting overcomes the limitations of

methods that require topic-specific demonstrations, thereby enabling effective cross-domain framing analysis. Extensive evaluations across multiple topics demonstrate that CWOS not only articulates a substantial portion of known frames but also uncovers many new frames that are both clearly articulated and sound.

## 8 Ethical Statement

We respected the privacy and honored the confidentiality of the authors who wrote the social media posts in the datasets from [Weinzierl and Harabagiu \(2022b\)](#) and [Mendelsohn et al. \(2021\)](#). We received approval from the Institutional Review Board at the University of Texas at Dallas for working with these Twitter / X social media datasets. IRB-21-515 stipulated that our research met the criteria for exemption #8(iii) of the Chapter 45 of Federal Regulations Part 46.101.(b). The experiments were conducted with a high level of professionalism, ensuring that the test collection was not used for evaluation until a final method was chosen based on initial results. All experimental details, including configurations and procedures, are thoroughly documented in the main text, supplementary materials, and the associated GitHub repository. We do not identify any significant risks associated with this research, which aims to enhance understanding of how controversial topics are discussed on social media. The research was guided by a strong commitment to the public good, with the overarching objective of advancing both natural language processing and communication research on social media.

## 9 Limitations

The approach presented in this paper, designed to discover and articulate Frames of Communication, is specifically tailored to social media content from Twitter / X. As a result, its effectiveness may be limited when applied to posts from other platforms, such as Reddit, where longer-form text is more common. Additionally, our current method relies solely on textual content, while many social media posts incorporate images, videos, and other forms of multimedia. In future work, we aim to enhance our methodology to account for the full multimodal nature of social media content, enabling frame discovery that incorporates visual and audiovisual elements. We also intend to adapt our approach for broader applicability across diverse social media

platforms.

A major limitation of the CWOS prompting approach is that it requires the use of LLMs capable of processing long contexts. The entire reference dataset must fit in the context, along with the annotated FoCs, as illustrated in Figure 4. Furthermore, there must be enough room in the context of the LLM to include the entire test context, as well as room to generate the structured response. This limitation explains why we were unable to include experiments with open-source models, as none at the time of performing experiments were capable of the required context lengths. However, we found one surprising consequence of this limitation: our API costs were drastically lower using CWOS prompting than with prior methods, such as CoT-ICACL. Though our method requires massive context sizes, CWOS only needs to run a single prompt through an LLM to produce a single generated response for the entire test dataset. Prior methods, such as CoT-ICACL, must prompt an LLM for every single post. Our method actually costs significantly less because of this difference.

## References

- Greg D. Adams. 1997. Abortion: Evidence of an issue evolution. *American journal of political science*, 41(3):718–737.
- Emily Allaway, Malavika Srikanth, and Kathleen McKeown. 2021. [Adversarial learning for zero-shot stance detection on social media](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4756–4767, Online. Association for Computational Linguistics.
- Sarah Banet-Weiser. 2018 - 2018. Empowered : popular feminism and popular misogyny.
- Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. [TweetEval: Unified benchmark and comparative evaluation for tweet classification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1644–1650, Online. Association for Computational Linguistics.
- STEVEN E. BARKAN. 2014. Gender and abortion attitudes: Religiosity as a suppressor variable. *Public opinion quarterly*, 78(4):940–950.
- Eric Baumer, Elisha Elovic, Ying Qin, Francesca Polletta, and Geri Gay. 2015. [Testing and comparing computational approaches for identifying the language of framing in political news](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics:*

- Human Language Technologies*, pages 1472–1482, Denver, Colorado. Association for Computational Linguistics.
- Rodney Benson. 2013a. *Shaping Immigration News: A French-American Comparison*. Communication, Society and Politics. Cambridge University Press.
- Rodney Benson. 2013b. *Shaping Immigration News: A French-American Comparison*. Communication, Society and Politics. Cambridge University Press.
- Francine D Blau and Lawrence M Kahn. 2017. The gender wage gap: Extent, trends, and explanations. *Journal of economic literature*, 55(3):789–865.
- Toby Bolsen and James N. Druckman. 2015. Counteracting the politicization of science. *Journal of communication*, 65(5):745–769.
- Toby Bolsen, James N. Druckman, and Fay Lomax Cook. 2014. How Frames Can Undermine Support for Scientific Adaptations: Politicization and the Status-Quo Bias. *Public Opinion Quarterly*, 78(1):1–26.
- Porismita Borah, Shreenita Ghosh, Jiyoun Suk, Darshana Sreedhar Mini, and Luhang Sun. 2023. Feminism not for all? the discourse around white feminism across five social media platforms. *Social media + society*, 9(3).
- Amber E. Boydston, Dallas Card, Justin Gross, Paul Resnick, and Noah A. Smith. 2018. [Tracking the development of media frames within and across policy issues](#).
- Dallas Card, Amber E. Boydston, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015a. [The media frames corpus: Annotations of frames across issues](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 438–444, Beijing, China. Association for Computational Linguistics.
- Dallas Card, Amber E. Boydston, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015b. The media frames corpus: Annotations of frames across issues. In *Annual Meeting of the Association for Computational Linguistics*.
- Dallas Card, Justin Gross, Amber Boydston, and Noah A. Smith. 2016. [Analyzing framing through the casts of characters in the news](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1410–1420, Austin, Texas. Association for Computational Linguistics.
- Dennis Chong and James N. Druckman. 2012. [Counter-framing effects](#). *The Journal of Politics*, 75(1):1–16.
- Elizabeth Adell Cook, Ted G Jelen, and Clyde Wilcox. 1993. Measuring public attitudes on abortion: Methodological and substantive considerations. *Perspectives on sexual and reproductive health*, 25(3):118–.
- Kimberle Crenshaw. 1991. Mapping the margins: Intersectionality, identity politics, and violence against women of color. *Stanford law review*, 43(6):1241–1299.
- Constance de Saint Laurent, Vlad Petre Glăveanu, and Claude Chaudet. 2020. Malevolent creativity and social media: Creating anti-immigration communities on twitter. *Creativity Research Journal*, 32:66 – 80.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- James N. Druckman and Mary C. McGrath. 2019. The evidence for motivated reasoning in climate change preference formation. *Nature climate change*, 9(2):111–119.
- Jiachen Du, Ruifeng Xu, Yulan He, and Lin Gui. 2017. [Stance classification with target-specific neural attention](#). In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 3988–3994.
- Robert M. Entman. 1993. Framing: Toward clarification of a fractured paradigm. *Journal of Communications*.
- Robert M. Entman. 2003. Cascading activation: Contesting the white house’s frame after 9/11. *Political Communication*, 20:415 – 432.
- Susan. Faludi. 1991. Backlash : the undeclared war against american women.
- MM Ferree. 2003. Resonance and radicalism: Feminist framing in the abortion debates of the united states and germany. *The American journal of sociology*, 109(2):304–344.
- Anjalie Field, Doron Kliger, Shuly Wintner, Jennifer Pan, Dan Jurafsky, and Yulia Tsvetkov. 2018. [Framing and agenda-setting in russian news: a computational analysis of intricate political strategies](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Diana Greene Foster. 2020 - 2020. The turnaway study : ten years, a thousand women, and the consequences of having—or being denied—an abortion.
- William A. Gamson. 1989. News as framing: Comments on graber. *The American Behavioral Scientist*, 33(2):157–161.
- Mattis Geiger, Franziska Rees, Lau Lilleholt, Ana P. Santana, Ingo Zettler, Oliver Wilhelm, Cornelia Betsch, and Robert Böhm. 2021. Measuring the 7cs of vaccination readiness. *European Journal of Psychological Assessment*, pages 1–9.

- Catherine Hakim. 2006. Women, careers, and work-life preferences. *British journal of guidance & counselling*, 34(3):279–294.
- Aífe Hopkins-Doyle et al. 2024. The misandry myth: An inaccurate stereotype about feminists’ attitudes toward men. *Psychology of women quarterly*, 48(1):8–37.
- Jan Fredrik Hovden and Hilmar Mjelde. 2019a. Increasingly controversial, cultural, and political: The immigration debate in scandinavian newspapers 1970–2016. *Javnost - The Public*, 26(2):138–157.
- Jan Fredrik Hovden and Hilmar Mjelde. 2019b. Increasingly controversial, cultural, and political: The immigration debate in scandinavian newspapers 1970–2016. *Javnost - The Public*, 26(2):138–157.
- Tunazzina Islam and Dan Goldwasser. 2025. Uncovering latent arguments in social media messaging by employing LLMs-in-the-loop strategy. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 7397–7429, Albuquerque, New Mexico. Association for Computational Linguistics.
- Shanto Iyengar. 1990. Framing responsibility for political issues: The case of poverty. *Political Behavior*, 12(1):19–40.
- Shanto Iyengar. 1991. *Is anyone responsible? How television frames political issues*. American politics and political economy series. University of Chicago Press, Chicago, IL, US.
- Ted G. Jelen. 2014. The subjective bases of abortion attitudes: A cross national comparison of religious traditions. *Politics and religion*, 7(3):550–567.
- Janel Jett, Leigh Raymond, and Erin P. Hennes. 2024. Process skeptical populist framing of climate change in right-leaning media. *Environmental politics*, 33(5):799–819.
- Kirsti M. Jylhä and Kahl Hellmer. 2020. Right-wing populism and climate change denial: The roles of exclusionary and anti-egalitarian preferences, conservative ideology, and antiestablishment attitudes. *Analyses of social issues and public policy*, 20(1):315–335.
- Shima Khanehzar, Andrew Turpin, and Gosia Mikołajczak. 2019. Modeling political framing across policy issues and contexts. In *Australasian Language Technology Association Workshop*.
- Anna Kheyfets, Shubhecchha Dhaurali, Paige Feyock, Farinaz Khan, April Lockley, Brenna Miller, and Ndidiamaka Amutah-Onukagha. 2023. The impact of hostile abortion legislation on the united states maternal mortality crisis: a call for increased abortion education. *Frontiers in public health*, 11:1291668–.
- Haewoon Kwak, Jisun An, and Yong-Yeol Ahn. 2020a. A systematic media frame analysis of 1.5 million new york times articles from 2000 to 2017. *Proceedings of the 12th ACM Conference on Web Science*, pages 305–314.
- Haewoon Kwak, Jisun An, and Yong-Yeol Ahn. 2020b. A systematic media frame analysis of 1.5 million new york times articles from 2000 to 2017. In *Proceedings of the 12th ACM Conference on Web Science, WebSci ’20*, page 305–314, New York, NY, USA. Association for Computing Machinery.
- Anthony Leiserowitz. 2006. Climate change risk perception and policy preferences: The role of affect, imagery, and values. *Climatic change*, 77(1-2):45–72.
- Anthony Leiserowitz and Nicholas Smith. 2017. *Affective imagery, risk perceptions, and climate change communication*.
- Stephan Lewandowsky, Michael E. Mann, Nicholas J. L. Brown, and Harris Friedman. 2016. Science and the public: Debate, denial, and skepticism. *Journal of social and political psychology*, 4(2):537–553.
- Yingjie Li and Cornelia Caragea. 2019. Multi-task stance detection with sentiment and stance lexicons. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6299–6305, Hong Kong, China. Association for Computational Linguistics.
- Scott Liebertz and Jaclyn Bunch. 2021. Backfiring frames: Abortion politics, religion, and attitude resistance. *Politics and religion*, 14(3):403–430.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692.
- Matthew Lockwood. 2018. Right-wing populism and the climate change agenda: exploring the linkages. *Environmental politics*, 27(4):712–732.
- Kristin. Luker. 1984. Abortion and the politics of motherhood.
- Jörg Matthes and Matthias Kohring. 2008. The content analysis of media frames: Toward improving reliability and validity. *Journal of Communication*, 58(2):258–279.
- Christina Maxwell, Hema Preya Selvanathan, Sam Hames, Charlie R. Crimston, and Jolanda Jetten. 2025. A mixed-methods approach to understand victimization discourses by opposing feminist subgroups on social media. *British journal of social psychology*, 64(1):e12785–n/a.
- Dawn McCaffrey and Jennifer Keys. 2000. Competitive framing processes in the abortion debate: Polarization-vilification, frame saving, and frame debunking. *Sociological quarterly*, 41(1):41–61.

- Holly J. McCammon. 2022. A war of words over abortion: The legal-framing contest over the undue burden standard. *The Justice system journal*, 43(4):623–644.
- Aaron M. McCright and Riley E. Dunlap. 2011. The politicization of climate change and polarization in the american public’s views of global warming, 2001–2010. *Sociological quarterly*, 52(2):155–194.
- Mary L. McHugh. 2012. Interrater reliability: the kappa statistic. *Biochemia medica*, 22(3):276–282.
- Julia Mendelsohn, Ceren Budak, and David Jurgens. 2021. [Modeling framing in immigration discourse on social media](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2219–2263, Online. Association for Computational Linguistics.
- Sharon Meraz and Zizi Papacharissi. 2013. Networked gatekeeping and networked framing on #egypt. *The International Journal of Press/Politics*, 18:138 – 166.
- Sarah Miller, Laura R. Wherry, and Diana Greene Foster. 2023. The economic consequences of being denied an abortion. *American economic journal. Economic policy*, 15(1):394–437.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. [SemEval-2016 task 6: Detecting stance in tweets](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41, San Diego, California. Association for Computational Linguistics.
- Nona Naderi and Graeme Hirst. 2017. Classifying frames at the sentence level in news articles. In *Recent Advances in Natural Language Processing*.
- Jennifer C. Nash. 2008. re-thinking intersectionality. *Feminist review*, 89(89):1–15.
- W. Russell Neuman, Lauren Guggenheim, S. Mo Jang, and So Young Bae. 2014. The dynamics of public attention: Agenda-setting theory meets big data. *Journal of Communication*, 64:193–214.
- Abigail Newell. 2019. Pro-life and pro-woman: Complicating the anti-abortion narrative in america.
- Matthew C. Nisbet. 2009. Communicating climate change: Why frames matter for public engagement. *Environment : science and policy for sustainable development*, 51(2):12–23.
- OpenAI. 2023. [Gpt-4 technical report](#).
- OpenAI. 2025. [Openai o3-mini system card](#).
- OpenAI et al. 2024. [Openai o1 system card](#).
- Long Ouyang et al. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- Thomas E. Patterson. 1992. [Is anyone responsible? how television frames political issues](#). *American Political Science Review*, 86(4):1060–1061.
- Nina Perger. 2018. Anti-gender campaigns in europe: Mobilizing against equality. *AS. Andragoška spoznanja*, 24(3):95–96.
- Research Center Pew. 2024. Social media and news fact sheet. <https://www.pewresearch.org/journalism/fact-sheet/social-media-and-news-fact-sheet/>.
- Devin Proctor. 2022. The #tradwife persona and the rise of radicalized white domesticity. *Persona Studies*, 8(2):7–26.
- Stephen D. Reese. 2007. [The Framing Project: A Bridging Model for Media Research Revisited](#). *Journal of Communication*, 57(1):148–154.
- Stephen D. Reese, Oscar H. Gandy, and August E. (Eds.) Grant. 2001. [Framing Public Life: Perspectives on Media and Our Understanding of the Social World](#). Routledge.
- Loretta J. Ross and Rickie Solinger. 2017 - 2017. Reproductive justice : an introduction.
- Shamik Roy and Dan Goldwasser. 2020. Weakly supervised learning of nuanced frames for analyzing polarization in news media. In *Conference on Empirical Methods in Natural Language Processing*.
- Laurie A. Rudman and Peter Samuel. Glick. 2008. The social psychology of gender : how power and intimacy shape gender relations.
- W. Russell Neuman, Lauren Guggenheim, S. Mo Jang, and Soo Young Bae. 2014. [The Dynamics of Public Attention: Agenda-Setting Theory Meets Big Data](#). *Journal of Communication*, 64(2):193–214.
- Bertram Scheufele. 2004. [Framing-effects approach: A theoretical and methodological critique](#). *Communications*, 29(4):401–428.
- HOWARD SCHUMAN, STANLEY PRESSER, and JACOB LUDWIG. 1981. Context effects on survey responses to questions about abortion. *Public opinion quarterly*, 45(2):216–223.
- Reva B. Siegel. 2007. The new politics of abortion: An equality analysis of woman-protective abortion restrictions. *University of Illinois law review*, 2007(3):991–1053.
- Daniel Skinner. 2012. The politics of medical necessity in american abortion debates. *Politics & gender*, 8(1):1–24.

- Parinaz Sobhani, Saif Mohammad, and Svetlana Kiritchenko. 2016. [Detecting stance in tweets and analyzing its interaction with sentiment](#). In *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics*, pages 159–169, Berlin, Germany. Association for Computational Linguistics.
- Alexa Spence, Wouter Poortinga, and Nick Pidgeon. 2012. The psychological distance of climate change. *Risk analysis*, 32(6):957–972.
- Gemini Team et al. 2024a. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#).
- Gemini Team et al. 2024b. [Gemini: A family of highly capable multimodal models](#).
- PW Toller, EA Suter, and TC Trautman. 2004. Gender role identity and attitudes toward feminism. *Sex roles*, 51(1-2):85–90.
- Hugo Touvron et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).
- Imogen Tyler. 2007. The selfish feminist: Public images of women’s liberation. *Australian feminist studies*, 22(53):173–190.
- Sander van der Linden. 2015. The social-psychological determinants of climate change risk perceptions: Towards a comprehensive model. *Journal of environmental psychology*, 41:112–124.
- Sander van der Linden, Anthony Leiserowitz, Seth Rosenthal, and Edward Maibach. 2017. Inoculating the public against misinformation about climate change. *Global challenges*, 1(2).
- Baldwin Van Gorp. 2010. *Strategies to take subjectivity out of framing analysis*, pages 84–109. Routledge.
- Claes H. Vreese. 2005. [News framing: Theory and typology](#). *Information Design Journal*, 13(1):51–62.
- Dror Walter and Yotam Ophir. 2019. [News frame analysis: An inductive mixed-method computational approach](#). *Communication Methods and Measures*, 13(4):248–266.
- Elke U. Weber. 2010. What shapes perceptions of climate change? *Wiley interdisciplinary reviews. Climate change*, 1(3):332–342.
- Maxwell Weinzierl and Sanda Harabagiu. 2022a. [VaccineLies: A natural language resource for learning to recognize misinformation about the COVID-19 and HPV vaccines](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6967–6975, Marseille, France. European Language Resources Association.
- Maxwell Weinzierl and Sanda Harabagiu. 2024a. [Discovering and articulating frames of communication from social media using chain-of-thought reasoning](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1617–1631, St. Julian’s, Malta. Association for Computational Linguistics.
- Maxwell Weinzierl and Sanda Harabagiu. 2024b. [Tree-of-counterfactual prompting for zero-shot stance detection](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 861–880, Bangkok, Thailand. Association for Computational Linguistics.
- Maxwell A. Weinzierl and Sanda M. Harabagiu. 2021. [Automatic detection of covid-19 vaccine misinformation with graph link prediction](#). *Journal of Biomedical Informatics*, 124:103955.
- Maxwell A. Weinzierl and Sanda M. Harabagiu. 2022b. [From hesitancy framings to vaccine hesitancy profiles: A journey of stance, ontological commitments and moral foundations](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 16(1):1087–1097.
- Maxwell A. Weinzierl and Sanda M. Harabagiu. 2024c. [The impact of stance object type on the quality of stance detection](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 15942–15954, Torino, Italia. ELRA and ICCL.
- Yiqi Wu, Xiaodan Hu, Ziming Fu, Siling Zhou, and Jiangong Li. 2024. [Gpt-4o: Visual perception performance of multimodal large language models in piglet activity understanding](#).
- Chang Xu, Cécile Paris, Surya Nepal, Ross Sparks, Chong Long, and Yafang Wang. 2020. [DAN: dual-view representation learning for adapting stance classifiers to new domains](#). In *ECAI 2020 - 24th European Conference on Artificial Intelligence*, volume 325, pages 2260–2267.
- Songlin Yang, Roger Levy, and Yoon Kim. 2023. [Unsupervised discontinuous constituency parsing with mildly context-sensitive grammars](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5747–5766, Toronto, Canada. Association for Computational Linguistics.
- Jing Zheng, Jyh-Herng Chow, Zhongnan Shen, and Peng Xu. 2023. [Grammar-based decoding for improved compositional generalization in semantic parsing](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1399–1418, Toronto, Canada. Association for Computational Linguistics.

## A Controversial Problems

### A.1 Topic: COVID-19 Vaccines

Vaccine hesitancy was characterized by seven problems that increase or decrease an individual’s likelihood of getting vaccinated, as introduced by Geiger

<b>Problem</b>	<i>Definition</i>
<i>Confidence</i>	Trust in the security and effectiveness of vaccinations, the health authorities, and the health officials who recommend and develop vaccines.
<i>Complacency</i>	Complacency and laziness to get vaccinated due to low perceived risk of infections.
<i>Constraints</i>	Structural or psychological hurdles that make vaccination difficult or costly.
<i>Calculation</i>	Degree to which personal costs and benefits of vaccination are weighted.
<i>Collective Responsibility</i>	Willingness to protect others and to eliminate infectious diseases.
<i>Compliance</i>	Support for societal monitoring and sanctioning of people who are not vaccinated.
<i>Conspiracy</i>	Conspiracy thinking and belief in fake news related to vaccination.

Table 6: Controversial Problems associated with COVID-19 vaccine hesitancy.

et al. (2021). Each problem, along with its definition, is provided in Table 6. Prior work Weinzierl and Harabagiu (2024a,c) investigated the benefit of using this 7C model of vaccine hesitancy in framing COVID-19 vaccine hesitancy on social media, and found these problems and their definitions to be highly relevant to the COVID-19 vaccine hesitancy discourse.

## A.2 Topic: Immigration

The controversial problems surrounding the topic of immigration have been studied extensively (Patterson, 1992; Benson, 2013b; Hovden and Mjelde, 2019b; Mendelsohn et al., 2021). Table 7 lists the problems and their definitions included in IMM. These problems are informed by the Policy Frames Codebook, which provides a general-purpose way to structure and describe frame problems in political communication content (Boydston et al., 2018). Table 7 lists the problems identified in Mendelsohn et al. (2021), which combines the 14 issue-generic frame problems introduced by the Policy Frames Codebook (Boydston et al., 2018) and the Media Frames Corpus (Card et al., 2015a) with 11 issue-specific problems (Benson, 2013a; Hovden and Mjelde, 2019a) and 2 narrative-focused problems (Iyengar, 1990). The problems in Table 7 informed our framing analysis for IMM.

## A.3 Topic: Abortion

Discourse on abortion is entirely dominated by framing, influencing societal attitudes, policy decisions, and individual beliefs. The controversial framing problems outlined in Table 8 were identi-

fied through extensive research in political science, sociology, law, and media analysis, and were included in AB. Each problem is listed along with relevant literature, which shows how competing narratives - from personal rights to public health - shape the ideological battleground between pro-choice and pro-life perspectives.

The “*Women’s Rights*” and “*Fetal Rights*” problems encapsulate the fundamental conflict in abortion debates. On one side, proponents of abortion rights argue that women must have control over their bodies and reproductive futures (Newell, 2019; Luker, 1984; Siegel, 2007). In contrast, opponents emphasize that the fetus possesses inherent rights and personhood from conception. This dichotomy, which forms the core of the debate, has been extensively analyzed in studies that highlight how interpreting these opposing problems serves to mobilize distinct political constituencies.

The “*Sanctity of Life*” framing problem portrays abortion as morally unacceptable under any circumstances by invoking religious and ethical imperatives about the sacredness of human life (BARKAN, 2014; Jelen, 2014; Luker, 1984). In juxtaposition, the “*Personal Morality*” problem suggests that the decision to terminate a pregnancy is a nuanced ethical choice best left to individual conscience rather than imposed by the state. These moral framings are central to public discourse, as they not only define ideological boundaries, but also inform legal and policy debates.

Legal dimensions are another significant aspect of the abortion debate, captured by the “*Legality & Constitutional Rights*” and “*Legislative & States’ Rights*” problems. The former emphasizes abortion as a constitutionally protected right—grounded in privacy and liberty claims (Skinner, 2012; McCammon, 2022; Siegel, 2007) - while the latter argues that abortion policy should be determined by state legislatures in a democratic process (Adams, 1997; Cook et al., 1993). These problems illustrate how judicial and legislative arenas are battlegrounds for competing interpretations of rights and authority.

Health and safety issues also play a pivotal role in framing the debate. The “*Health & Safety*” problem highlights abortion as a safe, essential medical procedure when performed legally, whereas related problems such as “*Abortion Harms Women*” and “*Medical Necessity*” question whether abortion is ever medically indispensable or if it poses risks to women’s physical and psychological well-being (Kheyfets et al., 2023; Miller et al., 2023; Siegel,

<b>Problem</b>	<b>Description</b>
<i>Economic</i>	Financial implications of an issue.
<i>Capacity &amp; Resources</i>	The availability or lack of time, physical, human, or financial resources.
<i>Morality &amp; Ethics</i>	Perspectives compelled by religion or secular sense of ethics or social responsibility.
<i>Fairness &amp; Equality</i>	The (in)equality with which laws, punishments, rewards, and resources are distributed.
<i>Legality, Constitutionality &amp; Jurisdiction</i>	Court cases and existing laws that regulate policies; constitutional interpretation; legal processes such as seeking asylum or obtaining citizenship; jurisdiction.
<i>Crime &amp; Punishment</i>	The violation of policies in practice and the consequences of those violations.
<i>Security &amp; Defense</i>	Any threat to a person, group, or nation and defenses taken to avoid that threat.
<i>Health &amp; Safety</i>	Health and safety outcomes of a policy issue, discussions of health care.
<i>Quality of Life</i>	Effects on people's wealth, mobility, daily routines, community life, happiness, etc.
<i>Cultural Identity</i>	Social norms, trends, values, and customs; integration/assimilation efforts.
<i>Public Sentiment</i>	General social attitudes, protests, polling, interest groups, public passage of laws.
<i>Political Factors &amp; Implications</i>	Focus on politicians, political parties, governing bodies, political campaigns and debates; discussions of elections and voting.
<i>Policy Prescription &amp; Evaluation</i>	Discussions of existing or proposed policies and their effectiveness.
<i>External Regulation &amp; Reputation</i>	Relations between nations or states/provinces; agreements between governments; perceptions of one nation/state by another.
<i>Victim: Global Economy</i>	Immigrants are victims of global poverty, underdevelopment, and inequality.
<i>Victim: Humanitarian</i>	Immigrants experience economic, social, and political suffering and hardships.
<i>Victim: War</i>	Focus on war and violent conflict as reasons for immigration.
<i>Victim: Discrimination</i>	Immigrants are victims of racism, xenophobia, and religion-based discrimination.
<i>Hero: Cultural Diversity</i>	Highlights positive aspects of differences that immigrants bring to society.
<i>Hero: Integration</i>	Immigrants successfully adapt and fit into their host society.
<i>Hero: Worker</i>	Immigrants contribute to economic prosperity and are an important source of labor.
<i>Threat: Jobs</i>	Immigrants take nonimmigrants' jobs or lower their wages.
<i>Threat: Public Order</i>	Immigrants threaten public safety by breaking the law or spreading disease.
<i>Threat: Fiscal</i>	Immigrants abuse social service programs and are a burden on resources.
<i>Threat: National Cohesion</i>	Immigrants' cultural differences are a threat to national unity and social harmony.
<i>Episodic</i>	Message provides concrete information about specific people, places, or events.
<i>Thematic</i>	Message is more abstract, placing stories in broader political and social contexts.

Table 7: Descriptions of Controversial Problems interpreted by Frames of Communication in immigration discourse.

2007; Skinner, 2012). These conflicting narratives underscore the importance of framing in shaping public policy and clinical guidelines.

Economic considerations further add complexity to the debate. The “*Economic Consequences*” problem highlights how access to abortion is vital for women’s financial stability and career advancement, especially among low-income and marginalized groups (Miller et al., 2023; Foster, 2020 - 2020). In contrast, the “*Abortion as Exploitation*” problem contends that abortion practices may disproportionately target disadvantaged communities, thereby reinforcing narratives of racial and economic injustice (Ross and Solinger, 2017 - 2017; McCaffrey and Keys, 2000).

Social identity and cultural values are also deeply interwoven into the debate. The “*Gender Equality*” problem stresses that abortion access is fundamental for women to participate equally in education, employment, and public life. Conversely, the “*Motherhood & Traditional Roles*” problem argues that society should more robustly support

women who choose to carry their pregnancies to term, thereby reinforcing traditional family values (Siegel, 2007; Ross and Solinger, 2017 - 2017; Ferree, 2003). This discussion is extended by the “*Reproductive Justice*” problem, which broadens the debate to include issues of race, class, and comprehensive social support systems (Ross and Solinger, 2017 - 2017).

Finally, the debate is further enriched by how language itself is used as a framing tool. The “*Religious Freedom*” problem contends that abortion laws should either reflect religious values or maintain a strict separation between religion and state (Jelen, 2014; BARKAN, 2014). Meanwhile, the “*Public Opinion & Political Polarization*” problem demonstrates how ideological divides and partisan identities intensify the conflict over abortion (McCaffrey and Keys, 2000; Adams, 1997). The “*Language & Terminology*” problem emphasizes that the choice of words - such as “pro-choice” versus “pro-life” or “fetus” versus “unborn baby” - profoundly influences public perception and discourse

<b>Problem</b>	<b>Description</b>
<i>Women's Rights</i>	Emphasizes abortion as a fundamental right and an aspect of gender equality, framing the decision as a personal choice that the state should not interfere with. <a href="#">Newell (2019)</a> ; <a href="#">Luker (1984)</a> ; <a href="#">Siegel (2007)</a>
<i>Fetal Rights</i>	Frames abortion as the taking of a human life, prioritizing the moral and legal rights of the fetus, often invoking the concept of personhood from conception. <a href="#">Newell (2019)</a> ; <a href="#">Luker (1984)</a> ; <a href="#">Siegel (2007)</a>
<i>Sanctity of Life</i>	Portrays abortion as morally wrong under virtually all circumstances, often rooted in religious or ethical arguments regarding the sacredness of human life. <a href="#">BARKAN (2014)</a> ; <a href="#">Jelen (2014)</a> ; <a href="#">Luker (1984)</a>
<i>Personal Morality</i>	Emphasizes that abortion is a complex moral decision best left to individual conscience rather than dictated by the state. <a href="#">BARKAN (2014)</a> ; <a href="#">Jelen (2014)</a> ; <a href="#">Luker (1984)</a>
<i>Legality &amp; Constitutional Rights</i>	Focuses on abortion as a constitutional right tied to privacy and liberty, debating legal standards such as undue burden and personhood. <a href="#">Skinner (2012)</a> ; <a href="#">McCammon (2022)</a> ; <a href="#">Siegel (2007)</a>
<i>Legislative &amp; States' Rights</i>	Argues that abortion should be regulated by state legislatures rather than courts, emphasizing democratic decision-making and states' authority. <a href="#">Adams (1997)</a> ; <a href="#">Cook et al. (1993)</a> ; <a href="#">Liebertz and Bunch (2021)</a>
<i>Health &amp; Safety</i>	Positions abortion as a necessary medical procedure that is safe when legally performed, while highlighting the risks of illegal or unsafe procedures. <a href="#">Kheyfets et al. (2023)</a> ; <a href="#">Miller et al. (2023)</a>
<i>Abortion Harms Women</i>	Claims that abortion poses physical and psychological risks to women, using disputed evidence to support restrictions. <a href="#">Siegel (2007)</a> ; <a href="#">Skinner (2012)</a>
<i>Medical Necessity</i>	Debates whether abortion is ever truly necessary for maternal health, with pro-choice advocates highlighting exceptions and pro-life advocates arguing modern medicine diminishes the need. <a href="#">Siegel (2007)</a> ; <a href="#">Skinner (2012)</a>
<i>Economic Consequences</i>	Frames abortion as economically essential for women's financial stability and career prospects, especially affecting low-income and marginalized groups. <a href="#">Miller et al. (2023)</a> ; <a href="#">Foster (2020 - 2020)</a>
<i>Abortion as Exploitation</i>	Argues that abortion disproportionately harms disadvantaged communities, sometimes framed as a form of racial or economic oppression. <a href="#">Ross and Solinger (2017 - 2017)</a> ; <a href="#">McCaffrey and Keys (2000)</a>
<i>Gender Equality</i>	Asserts that access to abortion is essential for ensuring equal opportunities in education, employment, and public life. <a href="#">Ross and Solinger (2017 - 2017)</a> ; <a href="#">Siegel (2007)</a> ; <a href="#">Ferree (2003)</a>
<i>Motherhood &amp; Traditional Roles</i>	Frames abortion as undermining the value of motherhood, suggesting that women deserve more support to choose parenthood over termination. <a href="#">Siegel (2007)</a> ; <a href="#">Ross and Solinger (2017 - 2017)</a> ; <a href="#">Ferree (2003)</a>
<i>Reproductive Justice</i>	Expands the debate to include racial, economic, and social factors, emphasizing the right to parent in safe and supportive environments. <a href="#">Ross and Solinger (2017 - 2017)</a>
<i>Religious Freedom</i>	Argues that abortion laws should either reflect religious beliefs about life or avoid imposing a particular religious viewpoint, thus protecting individual religious freedom. <a href="#">Jelen (2014)</a> ; <a href="#">BARKAN (2014)</a>
<i>Public Opinion &amp; Political Polarization</i>	Examines how partisan divides and identity politics reinforce entrenched positions in the abortion debate, contributing to polarization. <a href="#">McCaffrey and Keys (2000)</a> ; <a href="#">Adams (1997)</a> ; <a href="#">Liebertz and Bunch (2021)</a>
<i>Language &amp; Terminology</i>	Highlights how strategic word choices (e.g., "pro-choice" vs. "pro-life", "fetus" vs. "unborn baby") frame the debate and influence public perception. <a href="#">Cook et al. (1993)</a> ; <a href="#">SCHUMAN et al. (1981)</a>

Table 8: Controversial Problems in the U.S. abortion discourse.

([Cook et al., 1993](#); [SCHUMAN et al., 1981](#)).

#### A.4 Topic: Climate Change

Framing plays a critical role in shaping public perceptions of climate change, influencing policy support, individual behaviors, and societal responses. The controversial framing problems in [Table 9](#) were identified through extensive research in climate communication, political psychology, and media studies, and were included in [CC](#). Each problem has been extensively documented in literature, demonstrating how specific narratives hinder effective engagement and policy action.

The "*Scientific Uncertainty & Climate Denial*"

problem capitalizes on the perceived lack of consensus in climate science, often amplified by misinformation campaigns ([Nisbet, 2009](#); [Lewandowsky et al., 2016](#); [van der Linden et al., 2017](#)). Studies show that media portrayal of uncertainty fosters public skepticism, delaying climate action and policy implementation. The "*Ideological Polarization & Identity*" problem embeds climate change within partisan divides, where beliefs about climate science are shaped by political ideology rather than evidence ([Druckman and McGrath, 2019](#); [Bolsen and Druckman, 2015](#)). Right-wing populist narratives, in particular, have framed climate action as a liberal or globalist agenda, reinforcing resistance

<b>Problem</b>	<b>Description</b>
<i>Scientific Uncertainty &amp; Climate Denial</i>	Framing climate science as uncertain or unsettled, casting doubt on the scientific consensus, or portraying climate change as a hoax. This reduces trust in experts and delays action. <a href="#">Nisbet (2009)</a> ; <a href="#">Lewandowsky et al. (2016)</a> ; <a href="#">van der Linden et al. (2017)</a>
<i>Ideological Polarization &amp; Identity</i>	Presenting climate change as a partisan issue aligned with a specific political ideology, leading to social division and resistance based on group identity rather than scientific facts. <a href="#">Druckman and McGrath (2019)</a> ; <a href="#">Bolsen and Druckman (2015)</a>
<i>Economic Trade-offs (Environment vs. Economy)</i>	Framing climate policies as harmful to economic growth, job security, and affordability, emphasizing short-term economic costs over long-term benefits or the costs of inaction. <a href="#">McCright and Dunlap (2011)</a> ; <a href="#">van der Linden (2015)</a>
<i>Elitist vs. Public Needs (Populist Framing)</i>	Portraying climate action as an agenda pushed by global elites, bureaucrats, and scientists, disconnected from the struggles of ordinary citizens. This frame suggests policies are unfair or imposed without grassroots support. <a href="#">Lockwood (2018)</a> ; <a href="#">Jylhä and Hellmer (2020)</a> ; <a href="#">Jett et al. (2024)</a>
<i>Catastrophic "Doomsday" Narratives &amp; Fatalism</i>	Overemphasizing apocalyptic consequences of climate change without solutions, leading to public fear, helplessness, and disengagement due to perceived inevitability of disaster. <a href="#">Leiserowitz (2006)</a> ; <a href="#">Leiserowitz and Smith (2017)</a>
<i>Psychological Distance &amp; Abstraction</i>	Framing climate change as a distant, abstract issue affecting future generations, remote places, or complex scientific processes, making it seem less relevant to people's immediate lives and concerns. <a href="#">Weber (2010)</a> ; <a href="#">Spence et al. (2012)</a>

Table 9: Controversial Problems in climate change discourse.

to mitigation policies.

The “*Economic Trade-offs (Environment vs. Economy)*” problem presents climate action as detrimental to economic stability, emphasizing immediate costs over long-term benefits ([McCright and Dunlap, 2011](#); [van der Linden, 2015](#)). This problem has been shown to reduce public support for carbon pricing, renewable energy investment, and regulatory policies, as financial concerns dominate public discourse. The “*Elitist vs. Public Needs (Populist Framing)*” narrative portrays climate action as a top-down initiative imposed by distant elites, detached from the realities of ordinary citizens [Lockwood \(2018\)](#); [Jylhä and Hellmer \(2020\)](#); [Jett et al. \(2024\)](#). Populist movements often use this problem to undermine trust in climate science and governance, framing climate policies as unjust or unnecessary.

The “*Catastrophic "Doomsday" Narratives & Fatalism*” problem stems from excessive use of apocalyptic messaging, which can backfire by fostering helplessness and disengagement [Leiserowitz \(2006\)](#); [Leiserowitz and Smith \(2017\)](#). While highlighting risks is crucial, research suggests that overly dire predictions without actionable solutions lead to fatalistic attitudes rather than proactive behavior. Lastly, “*Psychological Distance & Abstraction*” refers to framing climate change as a distant problem - either temporally (impacting future generations), spatially (affecting remote regions), or conceptually (scientifically complex) [Weber \(2010\)](#); [Spence et al. \(2012\)](#). This disconnect reduces the sense of personal relevance and urgency,

weakening public motivation to act.

## A.5 Topic: Feminism

Framing is a crucial factor in shaping public discourse on feminism, influencing societal attitudes, policy decisions, and individual beliefs. The controversial framing problems outlined in Table 10 were identified through extensive research in *gender studies*, *media analysis*, and *political communication*, and were included in **FM**. These problems define how feminism is perceived, supported, or opposed in contemporary social and political debates. Each framing problem is rooted in the relevant literature, highlighting how different narratives shape feminist advocacy and resistance.

The “*Gender Equality vs. Misandry Myth*” problem captures the ongoing debate over whether feminism genuinely seeks gender equality or harbors hostility toward men. Supporters of feminism emphasize its role in dismantling gender-based oppression, whereas critics propagate the misandry myth, portraying feminism as an anti-male ideology [Ferree \(2003\)](#); [Rudman and Glick \(2008\)](#); [Borah et al. \(2023\)](#); [Hopkins-Doyle et al. \(2024\)](#). This problem is particularly prominent in digital discourse and media portrayals, often serving to delegitimize feminist activism.

The “*Victimhood vs. Empowerment*” problem questions whether feminism empowers women by addressing systemic inequalities or fosters a “victim mentality.” Feminist scholars argue that recognizing oppression is the first step toward dismantling discriminatory structures, while opponents claim that modern feminism overempha-

<b>Problem</b>	<b>Description</b>
<i>Gender Equality vs. Misandry Myth</i>	The debate over whether feminism promotes gender equality for all or fosters hostility toward men. Supporters argue feminism seeks to end gender-based oppression, while critics claim it is anti-male and seeks female superiority. Ferree (2003); Rudman and Glick (2008); Borah et al. (2023); Hopkins-Doyle et al. (2024)
<i>Victimhood vs. Empowerment</i>	The question of whether feminism empowers women by addressing systemic inequality or fosters a 'victim mentality.' Supporters argue that acknowledging discrimination is a step toward empowerment, while critics claim it discourages self-reliance and overemphasizes oppression. Banet-Weiser (2018 - 2018); Rudman and Glick (2008); Maxwell et al. (2025)
<i>Traditional Gender Roles vs. Feminist Ideals</i>	The conflict between feminism and traditional gender roles. Feminists argue for freedom of choice in career and family life, while critics claim feminism devalues motherhood, homemaking, and femininity. Toller et al. (2004); Proctor (2022)
<i>Intersectionality vs. White Feminism</i>	The tension between feminism's goal of inclusivity across race, class, and sexuality versus critiques that mainstream feminism is dominated by white, Western, and middle-class perspectives, often ignoring marginalized groups. Crenshaw (1991); Nash (2008); Borah et al. (2023)
<i>Economic Equality and the Pay Gap Debate</i>	The dispute over whether the gender wage gap is a result of discrimination or personal choices. Feminists argue that structural barriers cause persistent inequality, while critics claim the gap is exaggerated and primarily driven by women's career and lifestyle decisions. Blau and Kahn (2017); Hakim (2006)
<i>Policy Debates and Political Polarization</i>	The framing of feminist policies as either necessary for gender justice or as ideological overreach. Feminist goals such as reproductive rights and gender quotas are framed by supporters as essential for equality, while critics see them as attacks on traditional values or political extremism. Perger (2018); Ferree (2003)
<i>Media Representation and Backlash</i>	The portrayal of feminism and feminists in the media, where supporters argue for accurate representation of feminist goals, while detractors frame feminists as extremists, humorless, or irrelevant. This problem includes media-driven backlash against feminist gains. Faludi (1991); Borah et al. (2023); Tyler (2007)

Table 10: Controversial Problems in modern feminism discourse.

sizes women's victimization, discouraging personal agency and resilience Banet-Weiser (2018 - 2018); Rudman and Glick (2008); Maxwell et al. (2025). This problem is frequently leveraged in discussions on workplace equity, sexual harassment, and legal protections for women.

The "*Traditional Gender Roles vs. Feminist Ideals*" debate centers on the perceived conflict between feminist principles and "conventional" gender roles such as wife, mother, and homemaker. Feminists advocate for women's autonomy in choosing career or family life without societal constraints, while critics argue that feminism devalues femininity and traditional family structures Toller et al. (2004); Proctor (2022). Research on the "tradwife" movement reveals framing this problem is actively used to counter feminist narratives by promoting traditional domestic roles.

The "*Intersectionality vs. White Feminism*" problem concerns whether feminism is sufficiently inclusive of diverse racial, economic, and sexual identities or remains predominantly focused on white, Western, middle-class women. Scholars have criticized mainstream feminism for marginalizing non-white and working-class women, leading to calls for an intersectional approach to feminist activism Crenshaw (1991); Nash (2008); Borah et al. (2023). Frames that interpret this problem

are especially relevant in contemporary feminist movements, where intersectionality is both championed and critiqued for being co-opted without substantive systemic change.

The "*Economic Equality and the Pay Gap Debate*" examines whether gender-based pay disparities are the result of structural discrimination or individual career choices. Feminists argue that persistent wage gaps stem from workplace bias and societal expectations that disadvantage women, while critics claim that differences in pay reflect occupational preferences and life decisions rather than discrimination Blau and Kahn (2017); Hakim (2006). This problem significantly influences public perceptions of equal pay laws, workplace diversity programs, and career advancement policies.

The "*Policy Debates and Political Polarization*" problem highlights the division between feminist policy goals - such as reproductive rights, gender quotas, and workplace protections - and the backlash framing these initiatives as ideological overreach. Feminist policies are often justified as necessary steps toward gender justice, while critics frame them as government overreach, political radicalism, or threats to traditional social structures Perger (2018); Ferree (2003). This problem is particularly visible in global political debates, where gender equality measures frequently become cen-

```

response_format:
  type: "json_schema"
  json_schema:
    name: frame_analysis
    schema:
      type: object
      properties:
        frames:
          description: The frames evoked by these posts.
          type: array
          items:
            type: object
            properties:
              problems:
                description: All of the names of the problems addressed by this frame of communication.
                type: array
                items:
                  type: enum
                  enum: {problem_names_list}
              frame:
                description: Articulate the evoked frame of communication as a single sentence.
                type: string
            required:
              - problems
              - frame
            additionalProperties: false
          required:
            - frames
          additionalProperties: false
      strict: true

```

Figure 5: The JSON constrained decoding schema for CWOS prompting, in YAML format.

tral to broader ideological conflicts.

The “*Media Representation and Backlash*” problem explores the portrayal of feminists in mainstream and digital media, as well as the recurring backlash against feminist movements. While feminist discourse has gained visibility in news and entertainment media, anti-feminist backlash remains prevalent, often depicting feminists as aggressive, humorless, or extreme Faludi (1991); Borah et al. (2023); Tyler (2007). Research highlights that backlash narratives frequently emerge in response to feminist progress, reinforcing harmful stereotypes and undermining feminist campaigns.

## B Prompting Details

Figures 5 and 6 illustrate the detailed schema and prompt used in Corpus-Wide One-Shot (CWOS) prompting for discovering and articulating FoCs. Figure 5 illustrates the exact structured schema used for constrained decoding of FoCs and the CPs they interpret. This schema ensures the LLMs generate an exact format in their response, ensuring our method produces only a list of FoCs and the problems they interpret. Figure 6 provides the full prompt provided to the LLM to instruct it to perform CWOS prompting for discovering and articulating FoCs directly from posts.

```

system_prompt: >-
You are an expert linguistic assistant tasked with performing a framing analysis on a dataset of social media posts.
Each post in the dataset addresses one or more "problems." When users on social media communicate, they articulate
"frames" to explain these problems, often by proposing explicit or implicit "causes." Your job is to identify the
problems each post addresses and articulate the frames of communication (a single sentence each) that convey how those
problems are being explained (i.e., the causes). You must produce output in JSON format, adhering strictly to the
provided structured schema. You will encounter controversial, biased, or misinformed frames in the posts; you must
capture these frames exactly, without paraphrasing away the cause or meaning. For example, "Vaccines are a tool to
control world population and institute a new world order" is correct, while "There is debate around the purpose of
vaccines" is incorrect. Avoid injecting your own or your creators' opinions, and do not soften or alter the users'
stated frames. You must produce a large number of distinct frames, more frames than problems, capturing many
perspectives. There must be at least 10% as many frames as there are total posts (for example, at least 75 frames for
750 posts). Aim to cover 70% to 80% of the posts, meaning at least 70% of them should evoke at least one frame. Each
problem should be reflected across multiple frames, addressing unique perspectives and episodic or thematic nuances of
discourse. Do not merge different posts into the same broad frame unless their content is nearly identical; prioritize
nuance and diversity of viewpoints. Each frame's single-sentence statement must be unique, and you must not reuse the
demonstration frames. Before finalizing your output, ensure you meet the 10% frames rule, cover at least 70% of posts,
provide each frame as a single sentence, and strictly follow the provided JSON schema. If these conditions are not met,
the output is invalid. Remember that the demonstration is for illustration only—your frames must reflect the actual
discourse of the dataset at hand.
user_prompt: |-
Topic: {topic}

Problems:
{problem_list}

Posts:
{post_list}

```

Figure 6: The full prompt for CWOS prompting.