

Discovery, Articulation and Analysis of Communication Framings

Anonymous ACL submission

Abstract

In social media discourse, the framing of a topic dictates the way opinions are formed. This paper presents two multi-agent frameworks: one for discovering *what* topic aspects are framed, and another for identifying *how* they are framed. This enables an analysis of a topic's discourse that reveals what framings are dominant while also providing an interpretation of the framings. Experiments across four different topics show promising results for discovering, articulating and analyzing framings, providing a novel understanding of controversial topics. ¹

1 Introduction

In social media discourse, a variety of issues are discussed and their *framing* dictates the way these issues are understood and how opinions are formed. Therefore, it is important to discover *what* issues are framed in various ways as well as *how* they are framed. This will allow us to understand *why* people form certain opinions about a given topic.

The notion of *framing* originates from communication theories (Entman, 2003; Reese et al., 2001; Scheufele, 2004; Chong and Druckman, 2012; Bolsen et al., 2014), where framing is believed to enable the portrayal of an issue from one perspective to the necessary exclusion of alternative perspectives, as noted in (Boydston and Gross, 2013). For example, when discussing the issue of the economic impact of immigration, if it is framed as immigrants contributing to the economy through their labor and consumption, it informs pro-immigration attitudes, excluding the perspective of stealing jobs from citizens, which in turn may lead to anti-immigration sentiment. Because nowadays a large portion of the population participates and forms opinions based on social media

¹Code: <https://anonymous.4open.science/r/Discovery-Articulation-Analysis-CFs-2845/>

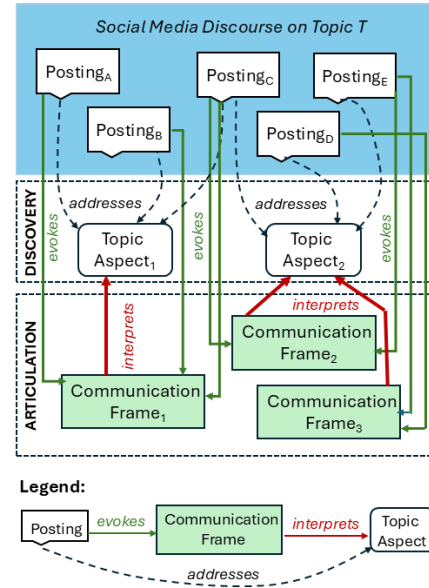


Figure 1: Social media postings addressing different topic aspects. The interpretation of the Topic Aspects (TAs) is provided by the articulation of the Communication Frames (CFs).

discourse, it is imperative to develop methods capable of capturing which issues are framed on social media, and especially how they are framed.

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.” Based on this definition, as illustrated in Figure 1, given the social media discourse on a topic T , revealing how the topic's issues are framed amounts to:

(a) the **discovery** of *what* Topic Aspects (TAs) are *addressed* in the various postings $\{P_A^T, P_B^T, P_C^T, P_D^T, P_E^T, \dots\}$, as shown in Figure 1. We note that the TAs correspond to the issues selected for framing in discourse, according to the definition from Entman (1993). For example, for the topic

056 $T = \text{immigration}$, we may have $TA_1 = \text{Economic}$
057 and $TA_2 = \text{Capacity and Resources}$. Appendix A
058 presents postings addressing TA_1 .

059 (b) the **articulation** of Communication
060 Frames (CFs) evoked by postings discussing
061 topic T . The CFs have the role of providing causal
062 interpretation to TAs. As shown in Figure 1, a CF
063 may interpret more than one TA, and more than
064 one CF may provide interpretations to a specific
065 TA. For example, both $TA_1 = \text{Economic}$ and
066 $TA_2 = \text{Capacity and Resources}$ are interpreted by
067 $CF_1 = \text{Immigrants largely spend their earnings}$
068 $\text{locally rather than sending all their money abroad,}$
069 $\text{so their wages are recycled into the local economy.}$
070 For each TA addressed by a posting P_i , the CF
071 evoked by P_i provides its causal interpretation.
072 The legend from Figure 1 summarizes the relations
073 established between TAs, CFs and postings.

074 (c) **analyzing** the topic framings by finding which
075 TAs are addressed the most. With millions of social
076 media postings produced daily, it is important
077 to have an understanding of what TAs are most
078 discussed in each topic discourse and to have an
079 immediate interpretations of these TAs through the
080 CFs that are evoked.

081 The problem of automatic TA discovery is not
082 new, nor is the problem of CF articulation. But
083 until now, TA discovery was cast as TA recog-
084 nition (Card et al., 2016; Naderi and Hirst, 2017;
085 Kwak et al., 2020), as it was assumed that for
086 each topic we are already aware of *all* TAs that are
087 typically framed in various ways when discussing
088 that topic. For example, Media Frames (Boydston
089 and Gross, 2013) proposed 14 TAs for topics in-
090 volving policy issues. These TAs emerged from a
091 combination of inductive and deductive strategies
092 painfully conducted by social scientists, as reported
093 in (de Vreese, 2005; Matthes and Kohring, 2008;
094 Walter and Ophir, 2019).

095 When considering the discovery of TAs for the
096 topic of immigration, (Mendelsohn et al., 2021)
097 used not only the 14 TAs introduced by the Media
098 Frames, but also 13 additional TAs which were spe-
099 cific to the topic of immigration, reported in studies
100 focusing on immigration discourse (Mahtani, 2016;
101 Hovden and Mjelde, 2019; Iyengar, 1991). This
102 shows that even when we are aware of some of
103 the TAs, we still should be able to *discover* addi-
104 tional TAs that help capturing all the ways a topic
105 is framed. This observation allowed us to formu-
106 late the research question: **RQ1: How could we**
107 *automatically discover all the TAs addressed in a*

108 *set of postings, when we known only some TAs?*

109 To answer RQ1, we developed **ID-TA**, a
110 multi-agentic LLM framework for Identifying and
111 Discovering TAs addressed in postings discussing a
112 certain topic. This multi-agent framework mimics
113 the deductive and inductive reasoning performed by
114 social scientists when trying to find and define the
115 TAs that are framed for a topic. To do so, **ID-TA**
116 is taking advantage of recent advances in reason-
117 ing with multi-agent LLMs through: (a) an *intro-*
118 *spective* refinement interaction between two LLM
119 agents (Sun et al., 2025); (b) *contrastive* reasoning
120 (Wu et al., 2024; Chia et al., 2023a; Jacovi et al.,
121 2021); and (c) multi-agent *debate* (Liang et al.,
122 2024). In this way, **ID-TA** is not only identify-
123 ing *known* TAs, but it also discovers new TAs and
124 validates them.

125 Although the TAs new and known of a topic can
126 indicate *which* topic framings are addressed in a
127 discourse, they do not account for the interpreta-
128 tions of the framings, which indicate *why* certain
129 TAs are addressed. This observation allowed us to
130 formulate the research question: **RQ2: How could**
131 *we automatically articulate the interpretations of*
132 *the TAs discovered in a topic discourse?*

133 To answer RQ2, we developed **Art-FRAME**,
134 a multi-agentic LLM framework capable to (1) ar-
135 ticulate FCs evoked by posting that addressed cer-
136 tain TAs, (2) filter out CFs that paraphrase already
137 articulated CFs; (3) validate the CFs through an ex-
138 plainability test. Although **Art-FRAME** takes
139 advantage of recently reported methods relying on
140 LLM reasoning (Weinzierl and Harabagiu, 2025,
141 2024), it considers a new validation of CFs.

142 This paper makes three contributions: (1) it
143 presents **ID-TA**, the first method capable to dis-
144 cover all topic aspects discussed in a discourse,
145 known or new, across several topics; (2) it presents
146 **Art-FRAME**, a method capable to articulate all
147 valid CFs evoked in the discourse of several topics;
148 and (3) it reports a novel analysis of controversial
149 topics, focused on the way topics are framed.

150 The remainder of the paper is organized as fol-
151 lows. Section 2 is detailing the methodology that al-
152 lows to discover framings and articulate them. Sec-
153 tion 3 details the data used in experiments across
154 discourse focused on five different topics. Section 5
155 describes the results of our experiments while Sec-
156 tion 5 showcases framing analysis of a topic’s dis-
157 course. Section 6 summarizes the conclusions.

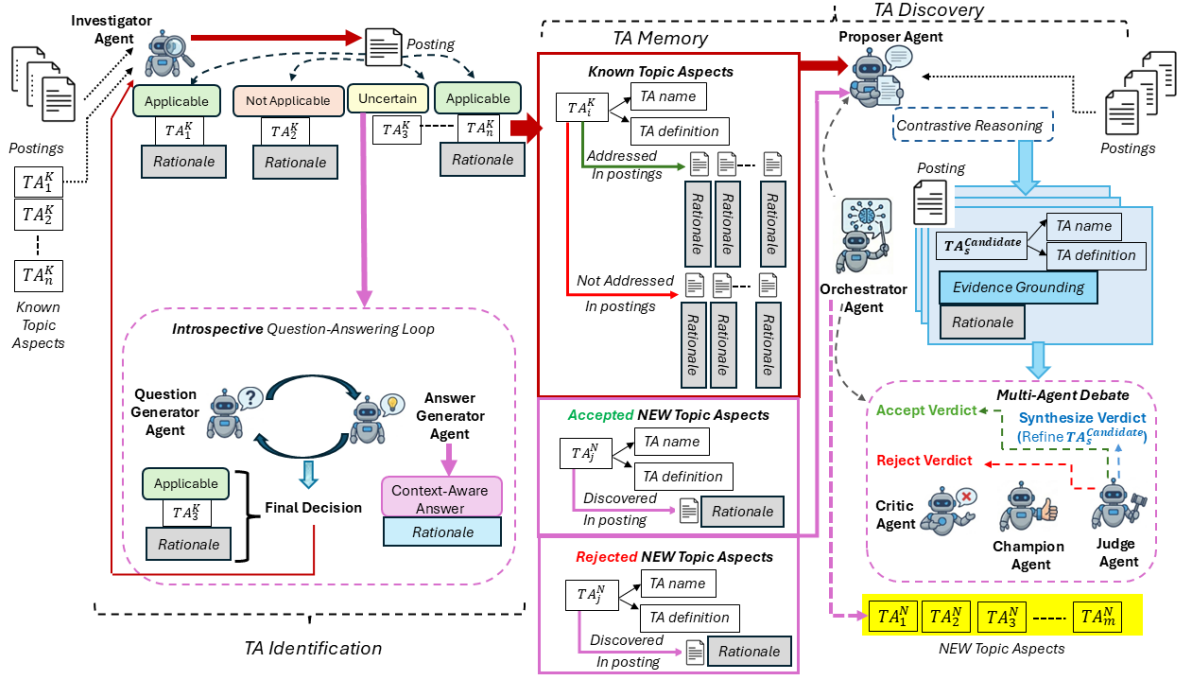


Figure 2: ID-TA: A Multi-Agent Module for Topic Aspect Discovery.

2 The Method

We propose a framework for processing framings in social media discourse consisting of two modules operating in sequential mode:

Module 1 is using the **ID-TA** multi-agent framework, described in Section 2.1

Module 2 is using the **Art-FRAME** multi-agent framework, described in Section 2.2

2.1 Module 1: ID-TA

As illustrated in Figure 2, before discovering new TAs from a set of postings discussing a topic T , **ID-TA** first addresses the problem of *identifying* known TAs in the postings. Appendix B.1 lists the known TAs considered for four different topics. **ID-TA** uses a *TA memory* that keeps track of which TAs are addressed in which postings. In addition, the TA memory informs the discovery of new TAs.

Identifying known TAs: An Investigator Agent is used to decide for each posting P_i which of the known TAs, $\{TA_1^K, TA_2^K, \dots, TA_n^K\}$ are addressed in P_i . The **Investigator Agent**, using the prompt detailed in Appendix C, can decide for any TA_x^K whether it is:

- applicable*, because P_i is addressing TA_x^K ; or
- not applicable*, because P_i is not addressing TA_x^K ; or
- uncertain*, when it does not know if P_i is addressing TA_x^K .

When the **Investigator Agent** is uncertain about a TA_x^K , it triggers the *Introspective Question-Answering Loop* shown in Figure 2, which resolves ambiguity by simulating a dialogue between two specialized sub-agents: a **Question Generator Agent** and an **Answer Generator Agent**.

The **Question Generator Agent** formulates a specific *Yes / No* question designed to test the applicability of TA_x^K . The **Answer Generator Agent** carefully analyses the post P_i to answer the question. It then proceeds to provide a definitive *Yes* or *No* answer to the specific question, accompanied by a rationale explaining the decision. If the **Answer Generator Agent** cannot determine an answer for a question, it requests the **Question Generator Agent** to formulate a different or more granular question. This Q/A loop between the two agents continues until either a definitive answer is found or the threshold \max_loop is reached.

The final decision of the **Investigator Agent** regarding the applicability of TA_x^K to P_i is resolved by inspecting the Question-Answer (Q/A) pairs. When the Investigator Agent finds sufficient evidence in the QA pairs that either the TA_x^K is applicable to P_i , or that TA_x^K is not applicable to P_i , it will change the uncertain decision to one of these decisions and it will also generate a *rationale* for the decision. If no evidence is found, it will decide that TA_x^K is not applicable to P_i . Detailed examples illustrating the introspective question-

answering loop are presented in Appendix D.

The TA Memory: Figure 2 shows how for each known TA_i^K , the memory keeps track of (a) the TA name (e.g. *Economy*); (b) its definition (e.g. *The costs, benefits, or monetary/financial implications*); (c) the list of the postings that addressed the TA; and eventually (d) the list of all postings that do not address the TA. The TA Memory is essential to the discovery of new TAs, and will be updated when new TAs are revealed.

Discovery of new TAs: is performed under the supervision of a central LLM **Orchestrator Agent**, which governs the life-cycle of new TA discovery through its three steps:

Step 1: For each posting P_j , the **Orchestrator Agent** invokes a **Proposer Agent** to generate a list of candidate new TAs $\{TA_i^{Candidate}\}$ by using *Contrastive Chain-of-Thought (CCoT)* reasoning (Chia et al., 2023b). Prior work has shown that such contrastive formulations improve reasoning quality and conceptual clarity (Jacovi et al., 2021; Chia et al., 2023b; Miller, 2019). Unlike simple Chain-of-Thought (CoT) reasoning (Wei et al., 2022), CCoT considers not only positive examples of demonstrations, but also negative examples. These examples are retrieved by the **Proposer Agent** from the TA Memory, where for each TA, examples of postings addressing it as well as examples of postings that do not address it are readily available. The CCoT prompt is detailed in Appendix C. For each $TA_i^{Candidate}$ generated by the CCoT prompt of the **Proposer Agent**, we also obtain: (a) the name of the $TA_i^{Candidate}$, e.g. *Victim Scapegoating*; (b) the definition of $TA_i^{Candidate}$, e.g. *Assigning unrelated social or policy failures to immigrants*; (c) evidence from P_j justifying that $TA_i^{Candidate}$ is being addressed; and (d) the rationale.

Step 2: The **Orchestrator Agent** checks if any of the $TA_i^{Candidate}$ are already known. For this reason, semantic similarity between the definitions of $TA_i^{Candidate}$ and any (known or new) $TA_j^{K/N}$ encoded in the TA memory is computed. We have considered this similarity between definitions $S_{i,j}^{Def}(TA_i^{Candidate}, TA_j^{K/N})$ to be measured by: (a) Sentence-BERT (SBERT) (Reimers and Gurevych, 2019) (b) BertScore (Zhang* et al., 2020); and (c) the Cross-Encoder introduced by (Nogueira and Cho, 2020). To decide which similarity measure to use, we have also relied on human comparisons of the pairs of definitions of TAs,

judged with a three-point Likert scale: Duplicate (0), Overlap (0.5), or Distinct (1.0). When cross-validating the human judgments with the results of each considered similarity metric, we have found that SBERT produced the best scores. If $S_{i,j}^{Def} \geq S_K$, where S_K is a similarity threshold, then the **Orchestrator Agent** considers $TA_i^{Candidate}$ to be a duplicate of a TA already encoded in the TA Memory. Otherwise, the **Orchestrator Agent** proceeds to its third step.

Step 3: The **Orchestrator Agent** invokes a multi-agent debate framework where the *discriminant validity* and *applicability* of a TA_{cand} are debated. Three different agents, namely a **Champion Agent**, as **Critic Agent** and a **Judge Agent** participate in the debate. During the debate over TA_i^{cand} , associated with a posting P_j , the agents roles are the following:

- ◇ **The Champion Agent** defends TA_i^{cand} , arguing for (a) its *applicability* based on its grounding in posting P_j ; and (b) the *discriminant validity*, based on TA_i^{cand} distinctness from all the known TAs addressed by P_j .

- ◇ **The Critic Agent** acts as a strict gatekeeper for the necessity of TA_i^{cand} . This agent operates on the principle of parsimony: new TAs should strictly not be created if existing ones suffice. It attempts to invalidate TA_i^{cand} by scrutinizing the **Champion Agent**'s arguments for (a) redundancy (when demonstrating TA_{cand} is similar to other TAs addressed in P_j); and (b) hallucination (demonstrating TA_{cand} relies on overreaching assumptions).

- ◇ **The Judge Agent** serves as the sole decision-maker and enforces a *dynamic termination policy*. Its verdict may take one of three forms: *Accept*, *Reject*, or *Correction*. When the arguments presented by any of the other two agents are decisive, the **Judge Agent** issues an early verdict to avoid unnecessary computation. Crucially, the **Judge Agent** also possesses *correction capabilities*: if the **Champion Agent** identifies a misalignment between the name or definitions of the TA encoded in the TA Memory and the candidate TA, the **Judge Agent** intervenes to refine the name or the definition of TA_i^{cand} . A detailed example of the multi-agent debate where corrections are performed is shown in Appendix E. The Appendix also shows the prompts used by the **Champion, Critic and Judge Agents**.

The advantage of using three different agents in this debate stems from the observation that in step 1, when the Proposer Agent utilizes contrastive reasoning to hypothesize a TA_i^{cand} for a posting

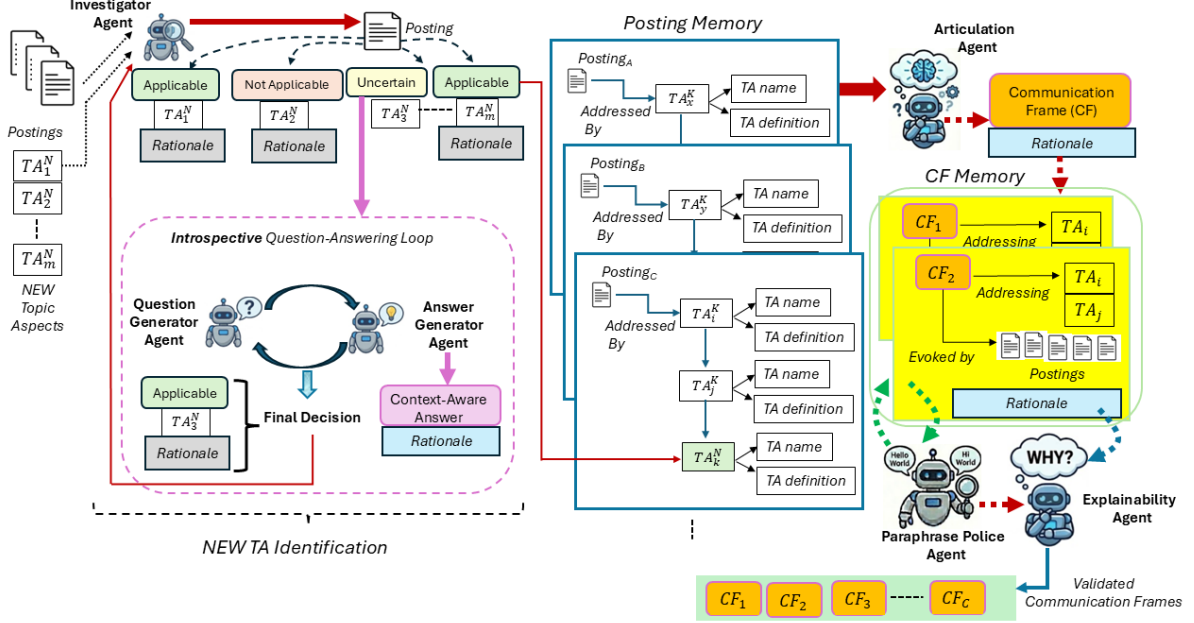


Figure 3: **Art-Frame**: A Multi-Agent Module for Articulating Communication Frames.

P_j , it uses a single generative model susceptible to Confirmation Bias (Turpin et al., 2023; Sharma et al., 2025). To mitigate this, we treat the Proposer Agent’s output (P_{cand}) strictly as a Hypothesis, not a new TA. To validate each such hypothesis the Adversarial Multi-Agent Debate between three agents is employed. This form of debate is based on a paradigm rooted in Minsky’s *Society of Mind* theory (Minsky, 1986), which posits that intelligence emerges from the interaction of conflicting subprocesses. The Multi-Agent Debate is rigorously testing the hypothesis to ensure it is not a hallucination or a subtle variant of a existing TAs (Turpin et al., 2023; Liang et al., 2024; Xiong et al., 2023).

When the **Judge Agent** accepts TA_i^{cand} , a new entry in the TA Memory is created to encode it as a new TA. As shown in Figure 2, along with the TA name and definition, the posting P_j where it was discovered is also stored, as well as the rationale for its discovery. Similarly, if TA_i^{cand} is rejected, a similar encoding is created, indicating that TA_i^{cand} is not a new TA. After all postings are processed, **ID-TA** outputs the list of newly discovered TAs, as illustrated in Figure 2, linked to their structured information available from the TA Memory.

2.2 Module 2: **Art-FRAME**

As illustrated in Figure 3, in order to articulate the CFs evoked in each posting P_i , we need to know which TAs, known or new, are addressed in P_i . While **ID-TA** identifies the known TAs in each

posting, we need also in **Art-FRAME** to identify which new TAs are addressed in each posting, in the same way known TAs are identified in each posting by **ID-TA**. This allows **Art-FRAME** to assemble a *Posting Memory*, which keeps track of all TAs, known or new, addressed in each P_i . The Posting Memory informs an **Articulation Agent** that uses Chain-of-Thought (CoT) prompting (Wei et al., 2023) to generate, for each post P_i , a set of articulated Communication Frames $\{CF_i^k\}$. From the Posting Memory, the **Articulation Agent** is able to retrieve $ST_i = \{TA_i^j\}$, all TAs addressed by a posting P_i . ST_i along with P_i inform the prompting of the **Articulation Agent**. A $Rationale_i^j$ is generated for each articulated CF_i^k . All CFs generated in this way are evoked by P_i and provide a causal interpretation of the TAs from ST_i . Figure 3 shows how a *CF Memory* is created, to list all $\{TA_i^p\}$ (known or new) addressed by each articulated CF, along with all the postings that evoke it and the rationale for the CF.

The CF Memory, as illustrated in Figure 3, is used by two agents: A **Paraphrase-Police Agent**, whose role is to find paraphrasing CFs; and an **Explainability Agent**, whose role is to validate the articulated CFs. The **Paraphrase-Police Agent** identifies paraphrases between CFs that address the same TAs, using the same few-shot CoT prompting reported in (Ailneni and Harabagiu, 2025). When two CFs are paraphrases, one of them is randomly

379 selected and deleted from the CF Memory.

380 After the **Paraphrase-Police Agent** finishes to
381 consolidate the CF Memory, the **Explainability**
382 **Agent** starts operating. Drawing on prior research
383 that CFs must be articulated clearly to be effective
384 (Weinzierl and Harabagiu, 2024), we employ an
385 *Articulation Reward Model* R_{art} as a quality filter.
386 Implemented as a discriminative encoder, R_{art} is
387 fine-tuned using binary a cross-entropy loss on a
388 dataset comprising both valid and invalid CFs. Dur-
389 ing inference, the model predicts a scalar validity
390 score $s \in [0, 1]$ for each CF. A strict acceptance
391 threshold τ is used, ensuring that only CFs with
392 high articulation quality are considered.

393 In addition, we assume that valid CF_i generated
394 by the **Proposer Agent** should pass a *reconstruc-*
395 *tion test* performed by the **Explainability Agent**.
396 Given rationale R_i provided for each CF_i in the
397 CF Memory, as well as the TAs for which CF_i pro-
398 vides an interpretation, the **Explainability Agent**
399 is prompted to re-generate the CF_i^R . The CoT
400 prompt is detailed in Appendix C. If CF_i and CF_i^R
401 are found as being paraphrases by the **Paraphrase-**
402 **Police Agent**, then the reconstruction test is passed.
403 When the reconstruction test is not passed, CF_i
404 is not considered valid. As shown in Figure 3,
405 **Art-FRAME** outputs a list of valid CFs evoked
406 by all postings.

407 3 The Data

408 We use the data that was released publicly in
409 (Weinzierl and Harabagiu, 2025), which covers
410 four different discourse topics:

411 1. Dataset for Immigration (IMM): It consists of
412 2,730 posts from Twitter / X annotated with 14
413 issue-generic TAs introduced by the Policy Frames
414 Codebook (Boydston et al., 2018; Card et al.,
415 2015), 11 issue-specific TAs (Benson, 2013; Hov-
416 den and Mjelde, 2019), and 2 narrative-focused
417 TAs (Iyengar, 1990) surrounding the topic of
418 immigration. In this study, we focus only on the
419 14 issue-generic TAs.

420 2. Dataset for COVID-19 Vaccine Hesitancy (VH):
421 consists of the test split of the Twitter / X postings
422 released by COVAXFRAMES (Wei). This
423 dataset contains 2,815 postings annotated with 7
424 TAs and 113 CFs. These TAs correspond to seven
425 dimensions of vaccine hesitancy that influence
426 an individual’s likelihood of vaccination (Geiger
427 et al., 2022).

428 3. Dataset for Climate Change (CC) & 4. Dataset

429 for Abortion (AB): consist of 933 Twitter / X
430 postings about Abortion (AB) and 564 postings
431 about Climate Change (CC), originating from the
432 SemEval-2016 Task 6 dataset (Mohammad et al.,
433 2016). using the release from TweetEval (Barbieri
434 et al., 2020). The known TAs and their definitions
435 for both topics are provided in Appendix B.1.

436 4 Experimental Results

437 **Implementation Details** We used GPT-5-mini
438 both for **ID-TA** and **Art-FRAME** as the back-
439 bone LLM for all agents. In accordance with the
440 GPT-5 model family, all agents operate with a fixed
441 temperature of 1.0. For **ID-TA**, the introspective
442 question-answering loop is limited to (\max_loop
443 $= 2$). The multi-agent debate terminates either by
444 the decision of the **Judge Agent** or after (\max_run
445 $= 3$). In **Art-FRAME**, τ is set to 0.8, ensuring
446 that only CFs with high articulation quality are
447 validated.

448 **Quantitative Results:** The evaluation results
449 for **ID-TA** allowed us to quantify (1) the identifica-
450 tion of *known* TAs as well as (2) the discovery
451 of *new* TAs across all four topics we considered.
452 To evaluate the identified known TAs, we used the
453 annotations from the IMM dataset, the only dataset
454 with annotated known TAs and measured the Pre-
455 cision $P = 0.76$, Recall $R = 0.86$ and $F_1 = 0.80$
456 obtained by **ID-TA**. We note that on the IMM
457 dataset, (Mendelsohn et al., 2021) reported the
458 best F1-score of 0.657 for known TA identifica-
459 tion, which was outperformed by **ID-TA**.

460 **ID-TA** also discovers new TAs. For the evalu-
461 ation of the discovery of new TAs, two baselines
462 were considered:

463 \diamond (*Baseline 1*): TA Discovery through postings clus-
464 tering (TAD_{HAC}): Assuming that all postings that
465 address the same new TA belong to the same clus-
466 ter, we have used Hierarchical Agglomerative Clus-
467 tering (HAC) (Ward, 1963) with the single-link
468 strategy, measuring the similarity between post-
469 ings with the Sentence-BERT model (Reimers and
470 Gurevych, 2019). The name and the definition of
471 a new TA, corresponding to one of the obtained
472 clusters, was generated by an LLM (GPT-5-mini)
473 with the prompt detailed in Appendix C.

474 \diamond (*Baseline 2*) Zero-Shot Prompting (ZS-Prompt)
475 for TA Discovery: A Chain-of-Thought (CoT)
476 prompting, in a zero-shot manner, of an LLM (GPT-
477 5-mini) was used, conditioned on known TAs. The
478 number of new TAs produced by each baseline are

compared with the number of TAs discovered by **ID-TA** and listed in Table 1.

Method	IMM	CC	AB	VH
<i>TAD_{HAC}</i>	21	18	19	15
<i>ZS-Prompt</i>	34	27	29	22
ID-TA	13	10	11	9

Table 1: Number of new TAs discovered across four topics: Immigration (IMM), Climate Change (CC), Abortion (AB), and Vaccine Hesitancy (VH).

The evaluation of the multi-agent debate taking place in **ID-TA** measures the proportion of candidate TAs that are accepted, rejected, or corrected, illustrated in Table 2. Across all topics, a substantial proportion of candidate TAs proposed by the **Proposer Agent** are rejected by the Judge Agent, indicating that CCoT prompting alone is not sufficient for proposing new TAs. Among accepted TAs, a non-trivial fraction are corrected by the **Judge Agent**, indicating that adversarial debate often warrants refinement of the name or definition of new TAs.

Topic	Accepted (%)	Corrected (%)	Rejected (%)
IMM	43	9	48
CC	40	10	50
AB	42	8	50
VH	39	7	54

Table 2: Distribution of decisions of the **Judge Agent** during the multi-agent debate from ID-TA.

We also quantify the ability of a new TA to elicit CFs we compute the number of valid CFs articulated by **Art-FRAME** for each new TA, those discovered by **ID-TA** as well as those produced by the two baselines. Table 3 reports the average number of articulated CFs elicited per new TA for each topic. While baseline methods often discover a larger number of TAs, these TAs tend to support fewer causal interpretations. In contrast, TAs discovered by **ID-TA** consistently yield a higher number of validated CFs per TA, indicating that **ID-TA** discovers new TAs that are more interpretable.

Method	IMM	CC	AB	VH
<i>TAD_{HAC}</i>	5.6	4.0	2.7	6.8
<i>ZS-Prompt</i>	4.2	3.1	2.0	5.0
ID-TA	9.0	6.6	3.8	11.4

Table 3: Average number of Communication Frames (CFs) articulated per new TA across topics.

Criterion 1: Definition Clarity (DC) of new TAs: Human judges decide whether the TA definition is fluent and easy to understand and rate on a 5-point Likert scale, ranging from 1 (incomprehensible) to 5 (perfectly clear).

Criterion 2: Definition-Name Alignment (DNA): Judges are asked to judge the extent to which the definition explains, justifies, and faithfully reflects the meaning implied by the TA name. Judgments use a 5-point Likert scale, with a score of 1 indicates that the definition and name are semantically mismatched; while a score of 5 indicates that the definition offers a comprehensive explanation of the TA name.

Criterion 3: Posting Coverage (PC): Judges evaluate whether postings addressing a certain TA genuinely discuss it. Judges are shown the TA name and definition, along with five postings randomly sampled from the list of postings addressing that TA. For each posting, judges answer *Yes* if the posting clearly addresses the TA and *No*, otherwise. The PC score for a TA is computed as the proportion of postings receiving a *Yes* judgment over all postings addressing the same TA.

Criterion 4 Semantic Distinctiveness (SD): measures how clearly TAs differ in meaning. Sentence Bert (Reimers and Gurevych, 2019) is used to measure the similarity between the definitions of pairs of TAs. After selecting the most similar pairs of TAs, judges are asked to consider how distinctive the TAs are, based on their definitions, using a three-point scale: *Duplicate* (0), *Overlap* (0.5), or *Distinct* (1.0).

Table 4: Evaluation Criteria for Discovered New TAs.

Model	DC	DNA	PC	SD
ID-TA	4.30	4.29	0.94	0.88
<i>TAD_{HAC}</i>	3.10	3.35	0.48	0.65
<i>ZS-Prompt</i>	3.78	4.05	0.54	0.59

Table 5: Qualitative results for the discovery of New TAs for the topic of Immigration (IMM)

Qualitative Results are produced for (1) the discovered TAs; and (2) the articulated CFs.

Quality of the discovered TAs are revealed by human evaluation using four criteria, listed in Table 4, based on several Likert scales (Robinson, 2014). Table 5 reports the qualitative evaluation results for the IMM topic, whereas Appendix G.1 lists the results for all four topics. The results indicate that the TAs discovered by **ID-TA** produce evaluations which consistently outperform the results of the baselines across of the measures of all four criteria. The largest gains are observed in Post Addressing (PA) and Semantic Distinctiveness (SD), indicating that TAs discovered by **ID-TA** are not only well grounded in the underlying posts but also semantically non-redundant. Improvements in Definition Clarity (DC) and Definition-Name Alignment (DNA) further suggest that **ID-TA** produces interpretable and well-formed TA definitions.

Quality of the articulated CFs are revealed by

Framing Analysis of the topic of Immigration




Figure 4: Framing Analysis of the topic of Immigration.

two properties: (a) the *soundness* of the rationale provided when articulating a CF; (b) the *clarity* of the CF articulation. Two linguists were tasked to judge the soundness and clarity of CFs, with N_S being the number of CFs deemed sound, and N_C the number of CFs deemed clear. With N_T the total number of CFs produced by **Art-FRAME**, the *quality of reasoning* (Z) involved in articulating CFs is computed as $Z = N_S/N_T$ while the *quality of the articulation* (A) of CFs is $A = N_C/N_T$. The results of $Z \geq .97$ for and $A \geq .95$ across all topics indicates high quality of the articulated CFs (see Appendix G.2). Also high Inter-Annotator Agreement was observed (see Appendix H).

5 Framing Analysis

Figure 4 illustrates the framing analysis for the immigration (IMM) topic, showing (a) the number of postings that address both the known and the new TAs; (b) the number of CFs articulated for each

TA; and (c) examples of the most evoked frame for each TA. Appendix I details the analysis for the other three topics. For each topic, we can notice that the frame analysis allows us to understand which TAs are most discussed, enabling our understanding of *what* is aspects of a topic are framed the most. Moreover, the analysis shows the importance of discovering new TAs, as many of them are much more discussed than some of the known TAs. In addition, the CFs explain *why* some TAs are discussed the most and *how* they are framed.

6 Conclusion

This paper is describing **ID-TA**, an multi-agent framework for discovering topic aspects in social media discourse. It also details **Art-FRAME**, a second multi-agent framework capable to articulate the way topic aspects are framed. The operation in tandem of these two LLM architectures provides a novel way of understanding a topic's discourse.

7 Ethical Statement

We respected the privacy and honored the confidentiality of the authors who wrote the social media posts in the datasets from (Wei) and (Mendelsohn et al., 2021). We received approval from the Institutional Review Board at ANONYMIZED for working with these Twitter/X social media datasets, and no attempt was made to identify or profile individual users. All experimental settings, configurations, and procedures are clearly documented in this paper, the supplemental material, and the linked GitHub repository to support transparency and reproducibility. Human evaluations were conducted to assess the quality of the discovered Topic Aspects and the articulated Communication Frames, and to ensure the responsible and reliable use of human judgment, we measured and reported inter-annotator agreement (IAA). We do not identify any significant risks arising from this research. The goal of this work is to improve understanding of how public issues are discussed and framed in social media discourse, with a broader commitment to advancing research at the intersection of natural language processing and communication studies for the public good.

8 Limitations

This work has several limitations that point to directions for future research. First, our experiments are conducted exclusively on social media posts from a single platform, Twitter/X, and rely solely on textual content. In practice, many social media posts also incorporate images, videos, hyperlinks, and other forms of multimedia. In future work, we plan to extend our method to support multimodal data, which would enable its application to other social media platforms such as Instagram and Facebook, where discourse is often expressed through a combination of text, images, and videos. Second, the proposed framework assumes the availability of a predefined set of known Topic Aspects for each topic, which serves as the starting point for TA identification and discovery. As described in Section 2.1, the Proposer Agent relies on these known TAs to perform Contrastive Chain-of-Thought (CCoT) reasoning, and they are also required to ground the multi-agent debate. While ID-TA is designed to discover new TAs beyond this initial set, its performance is influenced by the coverage and quality of the known TA inventory. In domains where such prior knowledge is

unavailable, additional mechanisms may be needed to more effectively bootstrap the discovery process. Finally, although the multi-agent architecture is designed to mitigate common LLM failure modes such as hallucination and redundancy, the framework remains dependent on the reasoning capabilities of large language models. As a result, errors in interpretation or subtle model biases may propagate through the discovery and articulation pipeline. Moreover, the use of multi-agent debate and introspective question–answering loops increases computational cost and inference latency, which may limit scalability in large-scale or real-time settings.

References

- Greg D. Adams. 1997. Abortion: Evidence of an issue evolution. *American Journal of Political Science*, 41(3):718–737.
- Rakshitha Rao Ailneni and Sanda M. Harabagiu. 2025. [Automatically discovering how misogyny is framed on social media](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 12189–12208, Albuquerque, New Mexico. Association for Computational Linguistics.
- 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.
- Rodney Benson. 2013. *Shaping Immigration News: A French-American Comparison*. Communication, Society and Politics. Cambridge University Press.
- Toby Bolsen and James N. Druckman. 2015. Countering 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.
- 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](#). *Carnegie Mellon University*.

665	Amber E. Boydston and Justin H. Gross. 2013. Identifying media frames and frame dynamics within and across policy issues.	719
666		720
667		721
668	Dallas Card, Amber E. Boydston, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. The media frames corpus: Annotations of frames across issues. In <i>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)</i> , pages 438–444, Beijing, China. Association for Computational Linguistics.	722
669		723
670		724
671		725
672		726
673		727
674		728
675		729
676		730
677	Dallas Card, Justin Gross, Amber Boydston, and Noah A. Smith. 2016. Analyzing framing through the casts of characters in the news. In <i>Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing</i> , pages 1410–1420, Austin, Texas. Association for Computational Linguistics.	731
678		732
679		733
680		734
681		735
682		736
683	Yew Ken Chia, Guizhen Chen, Anh Tuan Luu, Soujanya Poria, and Lidong Bing. 2023a. Contrastive chain-of-thought prompting. <i>ArXiv</i> , abs/2311.09277.	737
684		738
685		739
686	Yew Ken Chia, Guizhen Chen, Luu Anh Tuan, Soujanya Poria, and Lidong Bing. 2023b. Contrastive chain-of-thought prompting. <i>Preprint</i> , arXiv:2311.09277.	740
687		741
688		742
689	Dennis Chong and James N. Druckman. 2012. Counterframing effects. <i>The Journal of Politics</i> , 75(1):1–16.	743
690		744
691	Elizabeth Adell Cook, Ted G. Jelen, and Clyde Wilcox. 1993. Measuring public attitudes on abortion: Methodological and substantive considerations. <i>Perspectives on Sexual and Reproductive Health</i> , 25(3):118–.	745
692		746
693		747
694		748
695		749
696	Claes H. de Vreese. 2005. News framing: Theory and typology. <i>Information Design Journal</i> , 13:51–62.	750
697		751
698	James N. Druckman and Mary C. McGrath. 2019. The evidence for motivated reasoning in climate change preference formation. <i>Nature Climate Change</i> , 9(2):111–119.	752
699		753
700		754
701		755
702	Robert M. Entman. 1993. Framing: Toward clarification of a fractured paradigm. <i>Journal of Communications</i> .	756
703		757
704	Robert M. Entman. 2003. Cascading activation: Contesting the white house’s frame after 9/11. <i>Political Communication</i> , 20:415 – 432.	758
705		759
706		760
707	Myra Marx Ferree. 2003. Resonance and radicalism: Feminist framing in the abortion debates of the united states and germany. <i>American Journal of Sociology</i> , 109(2):304–344.	761
708		762
709		763
710		764
711	Diana Greene Foster. 2020. <i>The Turnaway Study: Ten Years, a Thousand Women, and the Consequences of Having—or Being Denied—an Abortion.</i>	765
712		766
713		767
714	Mattis Geiger, Franziska Rees, Lau Lilleholt, Ana P. Santana, Ingo Zettler, Oliver Wilhelm, Cornelia Betsch, and Robert Böhm. 2022. Measuring the 7Cs of vaccination readiness. <i>European Journal of Psychological Assessment</i> , 38(4):261–269.	768
715		769
716		770
717		771
718		
	Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. Deberta: Decoding-enhanced bert with disentangled attention. In <i>International Conference on Learning Representations.</i>	
	Jan Fredrik Hovden and Hilmar Mjelde. 2019. Increasingly controversial, cultural, and political: The immigration debate in scandinavian newspapers 1970–2016. <i>Javnost - The Public</i> , 26(2):138–157.	
	Shanto Iyengar. 1990. Framing responsibility for political issues: The case of poverty. <i>Political Behavior</i> , 12(1):19–40.	
	Shanto Iyengar. 1991. Is anyone responsible? how television frames political issues. <i>Contemporary Sociology</i> , 21:467.	
	Alon Jacovi, Swabha Swayamdipta, Shauli Ravfogel, Yanai Elazar, Yejin Choi, and Yoav Goldberg. 2021. Contrastive explanations for model interpretability. In <i>Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing</i> , pages 1597–1611, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.	
	Ted G. Jelen. 2014. The subjective bases of abortion attitudes: A cross-national comparison of religious traditions. <i>Politics and Religion</i> , 7(3):550–567.	
	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. <i>Analyses of Social Issues and Public Policy</i> , 20(1):315–335.	
	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. <i>Frontiers in Public Health</i> , 11:1291668.	
	Klaus Krippendorff. 2011. Computing krippendorff’s alpha-reliability.	
	Haewoon Kwak, Jisun An, and Yong-Yeol Ahn. 2020. A systematic media frame analysis of 1.5 million new york times articles from 2000 to 2017. <i>Proceedings of the 12th ACM Conference on Web Science.</i>	
	Anthony Leiserowitz. 2006. Climate change risk perception and policy preferences: The role of affect, imagery, and values. <i>Climatic Change</i> , 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. <i>Journal of Social and Political Psychology</i> , 4(2):537–553.	

772	Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang,	Nona Naderi and Graeme Hirst. 2017. Classifying	826
773	Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and	frames at the sentence level in news articles . In <i>Re-</i>	827
774	Zhaopeng Tu. 2024. Encouraging divergent thinking	<i>cent Advances in Natural Language Processing</i> .	828
775	in large language models through multi-agent debate.		
776	In <i>Proceedings of the 2024 Conference on Empiri-</i>	Abigail Newell. 2019. Pro-life and pro-woman: Com-	829
777	<i>cal Methods in Natural Language Processing</i> , pages	plicating the anti-abortion narrative in america.	830
778	17889–17904, Miami, Florida, USA. Association for		
779	Computational Linguistics.	Matthew C. Nisbet. 2009. Communicating climate	831
		change: Why frames matter for public engagement.	832
780	Scott Liebertz and Jaclyn Bunch. 2021. Backfiring	<i>Environment: Science and Policy for Sustainable De-</i>	833
781	frames: Abortion politics, religion, and attitude resis-	<i>velopment</i> , 51(2):12–23.	834
782	tance. <i>Politics and Religion</i> , 14(3):403–430.		
		Rodrigo Nogueira and Kyunghyun Cho. 2020. Passage	835
783	Matthew Lockwood. 2018. Right-wing populism and	re-ranking with bert . <i>Preprint</i> , arXiv:1901.04085.	836
784	the climate change agenda: Exploring the linkages.		
785	<i>Environmental Politics</i> , 27(4):712–732.	Stephen D. Reese, Oscar H. Gandy, and August E. (Eds.)	837
		Grant. 2001. <i>Framing Public Life: Perspectives on</i>	838
786	Kristin Luker. 1984. <i>Abortion and the Politics of Moth-</i>	<i>Media and Our Understanding of the Social World</i> .	839
787	<i>erhood</i> .	Routledge.	840
		Nils Reimers and Iryna Gurevych. 2019. Sentence-	841
788	Minelle K. Mahtani. 2016. Shaping immigration news:	BERT: Sentence embeddings using Siamese BERT-	842
789	a french-american comparison . <i>Journal of Ethnic</i>	networks . In <i>Proceedings of the 2019 Conference on</i>	843
790	<i>and Migration Studies</i> , 42:680 – 683.	<i>Empirical Methods in Natural Language Processing</i>	844
		<i>and the 9th International Joint Conference on Natu-</i>	845
791	Jörg Matthes and Matthias Kohring. 2008. The content	<i>ral Language Processing (EMNLP-IJCNLP)</i> , pages	846
792	analysis of media frames: Toward improving reliabil-	3982–3992, Hong Kong, China. Association for Com-	847
793	ity and validity . <i>Journal of communication</i> . <i>Journal</i>	putational Linguistics.	848
794	<i>of Communication</i> , 58:258–279.		
		John Robinson. 2014. <i>Likert Scale</i> , pages 3620–3621.	849
795	Dawn McCaffrey and Jennifer Keys. 2000. Com-	Springer Netherlands, Dordrecht.	850
796	petitive framing processes in the abortion debate:		
797	Polarization-vilification, frame saving, and frame de-	Loretta J. Ross and Rickie Solinger. 2017. <i>Reproductive</i>	851
798	bunking. <i>Sociological Quarterly</i> , 41(1):41–61.	<i>Justice: An Introduction</i> .	852
		Bertram Scheufele. 2004. Framing-effects approach: A	853
799	Holly J. McCammon. 2022. A war of words over abor-	theoretical and methodological critique . <i>Communi-</i>	854
800	tion: The legal-framing contest over the undue bur-	<i>cations</i> , 29(4):401–428.	855
801	den standard. <i>Justice System Journal</i> , 43(4):623–		
802	644.	Howard Schuman, Stanley Presser, and Jacob Ludwig.	856
		1981. Context effects on survey responses to ques-	857
803	Aaron M. McCright and Riley E. Dunlap. 2011. The	tions about abortion. <i>Public Opinion Quarterly</i> ,	858
804	politicization of climate change and polarization	45(2):216–223.	859
805	in the american public’s views of global warming,		
806	2001–2010. <i>Sociological Quarterly</i> , 52(2):155–194.	Mrinank Sharma, Meg Tong, Tomasz Korbak, David	860
		Duvenaud, Amanda Askeell, Samuel R. Bow-	861
807	Julia Mendelsohn, Ceren Budak, and David Jurgens.	man, Newton Cheng, Esin Durmus, Zac Hatfield-	862
808	2021. Modeling framing in immigration discourse on	Dodds, Scott R. Johnston, Shauna Kravec, Timo-	863
809	social media. In <i>Proceedings of the 2021 Conference</i>	thy Maxwell, Sam McCandlish, Kamal Ndousse,	864
810	<i>of the North American Chapter of the Association</i>	Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda	865
811	<i>for Computational Linguistics: Human Language</i>	Zhang, and Ethan Perez. 2025. Towards under-	866
812	<i>Technologies</i> , pages 2219–2263, Online. Association	standing sycophancy in language models . <i>Preprint</i> ,	867
813	for Computational Linguistics.	arXiv:2310.13548.	868
		Reva B. Siegel. 2007. The new politics of abor-	869
814	Tim Miller. 2019. Explanation in artificial intelligence:	tion: An equality analysis of woman-protective abor-	870
815	Insights from the social sciences . <i>Artificial Intelli-</i>	tion restrictions. <i>University of Illinois Law Review</i> ,	871
816	<i>gence</i> , 267:1–38.	2007(3):991–1053.	872
		Daniel Skinner. 2012. The politics of medical necessity	873
817	Marvin Minsky. 1986. <i>The Society of Mind</i> . Simon &	in american abortion debates. <i>Politics & Gender</i> ,	874
818	Schuster.	8(1):1–24.	875
		Alexa Spence, Wouter Poortinga, and Nick Pidgeon.	876
819	Saif Mohammad, Svetlana Kiritchenko, Parinaz Sob-	2012. The psychological distance of climate change.	877
820	hani, Xiaodan Zhu, and Colin Cherry. 2016.	<i>Risk Analysis</i> , 32(6):957–972.	878
821	SemEval-2016 task 6: Detecting stance in tweets .		
822	In <i>Proceedings of the 10th International Workshop</i>		
823	<i>on Semantic Evaluation (SemEval-2016)</i> , pages 31–		
824	41, San Diego, California. Association for Computa-		
825	tional Linguistics.		

879	Li Sun, Liu He, Shuyue Jia, Yangfan He, and Chenyu You. 2025. DocAgent: An agentic framework for multi-modal long-context document understanding. In <i>Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing</i> , pages 17712–17727, Suzhou, China. Association for Computational Linguistics.	6814, Suzhou, China. Association for Computational Linguistics.	935
880			936
881			
882		Shirley Wu, Shiyu Zhao, Qian Huang, Kexin Huang, Michihiro Yasunaga, Vassilis N. Ioannidis, Karthik Subbian, Jure Leskovec, and James Zou. 2024. Avatar: Optimizing llm agents for tool usage via contrastive reasoning. <i>Advances in Neural Information Processing Systems 37</i> .	937
883			938
884			939
885			940
886	Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. 2023. Language models don't always say what they think: unfaithful explanations in chain-of-thought prompting. In <i>Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23</i> , Red Hook, NY, USA. Curran Associates Inc.		941
887			942
888		Kai Xiong, Xiao Ding, Yixin Cao, Ting Liu, and Bing Qin. 2023. Examining inter-consistency of large language models collaboration: An in-depth analysis via debate. In <i>Findings of the Association for Computational Linguistics: EMNLP 2023</i> , pages 7572–7590, Singapore. Association for Computational Linguistics.	943
889			944
890			945
891			946
892			947
893	Sander van der Linden. 2015. The social-psychological determinants of climate change risk perceptions: Towards a comprehensive model. <i>Journal of Environmental Psychology</i> , 41:112–124.		948
894			949
895		Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with bert. In <i>International Conference on Learning Representations</i> .	950
896			951
897	Sander van der Linden, Anthony Leiserowitz, Seth Rosenthal, and Edward Maibach. 2017. Inoculating the public against misinformation about climate change. <i>Global Challenges</i> , 1(2).		952
898			953
899			
900			
901	Dror Walter and Yotam Ophir. 2019. News frame analysis: An inductive mixed-method computational approach. <i>Communication Methods and Measures</i> , 13:248 – 266.	A Example of Postings, Topic Aspects and Communication Frames for Immigration	954
902			955
903			956
904			
905	Joe H. Ward. 1963. Hierarchical grouping to optimize an objective function. <i>Journal of the American Statistical Association</i> , 58:236–244.	Figure 5 illustrates an example of a Topic Aspect, namely TA_1 corresponding to the topic of immigration (IMM). The Figure also shows several postings discussing the topic and addressing TA_1 . The name and the definition of TA_1 are also provided.	957
906			958
907			959
908	Elke U. Weber. 2010. What shapes perceptions of climate change? <i>Wiley Interdisciplinary Reviews: Climate Change</i> , 1(3):332–342.		960
909			961
910			962
911	Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. Emergent abilities of large language models. <i>ArXiv</i> , abs/2206.07682.	The definition of TA_1 is also shown in Figure 5. Each TA has a Name and a Definition. As shown in Figure 5, the name of TA_1 is "Economic". The definitions of TA_1 explains its name. This is different from Communication Frames (CFs) which provide a causal interpretation to a TA. We note that each TA typically has multiple CFs that provide different interpretations.	963
912			964
913			965
914			966
915			967
916			968
917			969
918	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. Chain-of-thought prompting elicits reasoning in large language models. <i>Preprint</i> , arXiv:2201.11903.	When identifying TAs in postings, the definition of each TA informs the entailment of the TA from the posting. However, when new TAs need to be discovered, their name and definitions needs to be generated.	970
919			971
920			972
921			973
922			974
923	Maxwell Weinzierl and Sanda Harabagiu. 2024. Discovering and articulating frames of communication from social media using chain-of-thought reasoning. In <i>Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 1617–1631, St. Julian's, Malta. Association for Computational Linguistics.	The first two postings illustrated in Figure 5 evoke communication frame CF_1 , while the next two postings evoke communication frame CF_1 . Each of these communication frame provide a different causal interpretation of TA_1 .	975
924			976
925			977
926			978
927			979
928			980
929			981
930			982
931	Maxwell Weinzierl and Sanda M. Harabagiu. 2025. Investigating controversy framing across topics on social media. In <i>Findings of the Association for Computational Linguistics: EMNLP 2025</i> , pages 6795–	As shown in Figure 5, the first two postings evoke CF_1 which articulates a causal interpretation of TA_1 which (1) does not use any words appearing in the first two postings, is not a summary of those postings, but rather it expressed the impact of immigrants on economy as spending their earnings	983
932			984
933			985
934			

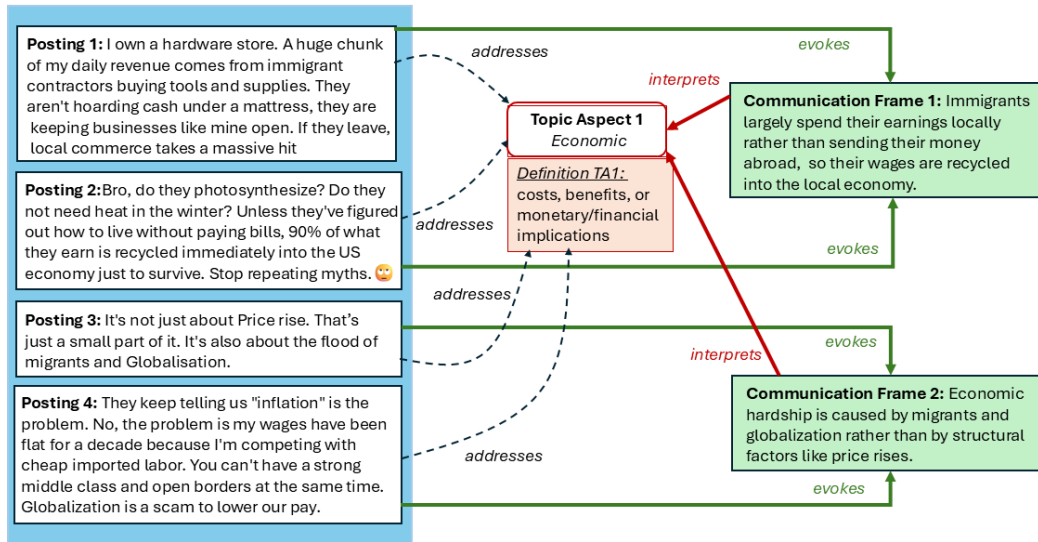


Figure 5: Example of an immigration Topic Aspect (TA) addressed by multiple postings and the Communication Frames (CFs) evoked by the postings.

locally instead of sending their money abroad, to their families.

The communication frame CF_2 illustrated in Figure 5 provides a different interpretation to TA_1 than the one offered by CF_1 , because it articulates the reasons for the economic hardship faced by some folks today, namely (1) the immigrants; and (2) globalization. It also states that price rises have a smaller effect. As with CF_1 , as shown in Figure 5, CF_2 is not a paraphrase or summary of the postings that evoke it.

In the analysis presented in this paper, we have found that TA_1 is interpreted by 17 different CFs. We have also found that this TA was one of the first four most addressed TAs for the topic of immigration. It was addressed by 591 different postings. However, just knowing that all the 591 postings address the $TA_1 = \text{Economic}$ aspect of immigration, does not inform us about the specific ways in which we need to interpret and understand TA_1 in order to comprehend and form an opinion on immigration, based on the discourse from social media. The communication frames CF_1 and CF_2 exemplify some of the 17 possible interpretations of TA_1 .

If one is aware of CF_1 , because they were exposed only to the first two postings illustrated in Figure 5, they might have a favorable opinion of immigration. In contrast, if someone else would have been exposed only to the last two postings illustrated in Figure 5, they would have inferred the communication frame CF_2 , and form an unfavor-

able opinion about immigration.

Knowing which communication frames are most evoked for each topic aspect can explain the dominant opinions about a topic discussed on social media. For example, if for topic aspect TA_1 shown in Figure 5 the dominant communication frame would be CF_2 , we could have an explanation for anti-immigrant opinions.

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

Table 6: Names and Definitions of Known Topic Aspects for the topic of Vaccine Hesitancy (VH).

B Topic Aspects and Definitions

B.1 Known Topic Aspects

Table 6 lists the known TAs and their definitions for the topic of COVID-19 Vaccine Hesitance

(VH), provided in (Weinzierl and Harabagiu, 2024). These TAs were informed by the seven factors reported in (Geiger et al., 2022) to increase or decrease an individual’s likelihood of getting vaccinated. The names and the definitions of these known TAs were provided in (Geiger et al., 2022).

Table 7 lists all the names and definitions for the known TAs for the topic of Immigration (IMM). These TAs represent the dimensions of framings introduced in (Card et al., 2015). Both the names and the definitions of these TAs have been reported in (Card et al., 2015).

For the topic of Climate Change (CC), Table 8 lists all the known TAs and their definitions. These TAs have been first described in (Weinzierl and Harabagiu, 2025). The TA names and their definitions originate, as reported in (Weinzierl and Harabagiu, 2025) is a multitude of academic papers addressing the topic of climate change, namely (Nisbet, 2009; Lewandowsky et al., 2016; van der Linden et al., 2017) for the first known TA; (Druckman and McGrath, 2019; Bolsen and Druckman, 2015) for the second known TA; (McCright and Dunlap, 2011; van der Linden, 2015) for the third known TA, (Lockwood, 2018; Jylhä and Hellmer, 2020) for the fourth known TA; (Leiserowitz, 2006; Leiserowitz and Smith, 2017) for the fifth known TA and (Weber, 2010; Spence et al., 2012) for the sixth known TA. This shows that an extensive inspection of academic literature is needed when requiring awareness of known TAs for some topics.

For the topic of Abortion (AB), we have relied on the known TAs and their definitions first reported in (Weinzierl and Harabagiu, 2025). The names and definitions of these 17 TAs are listed in Table 9. These TAs also resulted from the inspection of a large number of academic articles.

Table 9 lists of the TAs and their definitions for the topic of Abortion (AB), available from (Weinzierl and Harabagiu, 2025).

#. TA Name	TA Definition
1. <i>Economic</i>	The costs, benefits, or monetary/financial implications of immigration.
2. <i>Capacity & Resources</i>	The availability or lack of time, physical, human, or financial resources.
3. <i>Morality & Ethics</i>	Perspectives compelled by religion or secular sense of ethics or social responsibility.
4. <i>Fairness & Equality</i>	The (in)equality with which laws, punishments, rewards, resources are distributed.
5. <i>Legality, Constitutionality & Jurisdiction</i>	Court cases and existing laws that regulate policies; constitutional legal processes such as seeking asylum or obtaining citizenship; jurisdiction.
6. <i>Crime & Punishment</i>	The violation of policies in practice and the consequences of those violations.
7. <i>Security & Defense</i>	Any threat to a person, group, or nation and defenses taken to avoid that threat.
8. <i>Health & Safety</i>	Health and safety outcomes of a policy issue, discussions of health care.
9. <i>Quality of Life</i>	Effects on people’s wealth, mobility, daily routines, community life, happiness.
10. <i>Cultural Identity</i>	Social norms, trends, values, and customs; integration/assimilation efforts.
11. <i>Public Sentiment</i>	General social attitudes, protests, polling, interest groups, public passage of laws.
12. <i>Political Factors & Implications</i>	Focus on politicians, political parties, governing bodies, political campaigns and debates; discussions of elections and voting.
13. <i>Policy Prescription & Evaluation</i>	Discussions of existing or proposed policies and their effectiveness.
14. <i>External Regulation & Reputation</i>	Relations between nations or states/provinces; agreements between governments; perceptions of one nation/state by another.

Table 7: Names and Definitions of Known Topic Aspects for the topic of Immigration (IMM).

#. TA Name	TA Definition
1. <i>Scientific Uncertainty & 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. (Nisbet, 2009; Lewandowsky et al., 2016; van der Linden et al., 2017)
2. <i>Ideological Polarization & 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. (Druckman and McGrath, 2019; Bolsen and Druckman, 2015)
3. <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. (McCright and Dunlap, 2011; van der Linden, 2015)
4. <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. (Lockwood, 2018; Jylhä and Hellmer, 2020)
5. <i>Catastrophic “Doomsday” Narratives & Fatalism</i>	Overemphasizing apocalyptic consequences of climate change without solutions, leading to public fear, helplessness, and disengagement due to perceived inevitability of disaster. (Leiserowitz, 2006; Leiserowitz and Smith, 2017)
6. <i>Psychological Distance & 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. (Weber, 2010; Spence et al., 2012)

Table 8: Names and Definitions of Known Topic Aspects for Climate Change (CC).

#. TA Name	TA Definition
1. <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. (Newell, 2019; Luker, 1984; Siegel, 2007)
2. <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. (Newell, 2019; Luker, 1984; Siegel, 2007)
3. <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. (Barkan, 2014; Jelen, 2014; Luker, 1984)
4. <i>Personal Morality</i>	Emphasizes that abortion is a complex moral decision best left to individual conscience rather than dictated by the state. (Barkan, 2014; Jelen, 2014; Luker, 1984)
5. <i>Legality & Constitutional Rights</i>	Focuses on abortion as a constitutional right tied to privacy and liberty, debating legal standards such as undue burden and personhood. (Skinner, 2012; McCammon, 2022; Siegel, 2007)
6. <i>Legislative & States' Rights</i>	Argues that abortion should be regulated by state legislatures rather than courts, emphasizing democratic decision-making and states' authority. (Adams, 1997; Cook et al., 1993; Liebertz and Bunch, 2021)
7. <i>Health & Safety</i>	Positions abortion as a necessary medical procedure that is safe when legally performed, while highlighting the risks of illegal or unsafe procedures. (Kheyfets et al., 2023; Miller, 2019)
8. <i>Abortion Harms Women</i>	Claims that abortion poses physical and psychological risks to women, using disputed evidence to support restrictions. (Siegel, 2007; Skinner, 2012)
9. <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. (Siegel, 2007; Skinner, 2012)
10. <i>Economic Consequences</i>	Abortion is economically essential for women's financial stability and career prospects, especially affecting low-income and marginalized groups. (Miller, 2019; Foster, 2020)
11. <i>Abortion as Exploitation</i>	Argues that abortion disproportionately harms disadvantaged communities, sometimes framed as a form of racial or economic oppression. (Ross and Solinger, 2017; McCaffrey and Keys, 2000)
12. <i>Gender Equality</i>	Asserts that access to abortion is essential for ensuring equal opportunities in education, employment, and public life. (Ross and Solinger, 2017; Siegel, 2007; Ferree, 2003)
13. <i>Motherhood & Traditional Roles</i>	Abortion is undermining the value of motherhood, suggesting that women deserve more support to choose parenthood over termination. (Ross and Solinger, 2017; Siegel, 2007; Ferree, 2003)
14. <i>Reproductive Justice</i>	Expands the debate to include racial, economic, and social factors, emphasizing the right to parent in safe and supportive environments. (Ross and Solinger, 2017)
15. <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. (Jelen, 2014; Barkan, 2014)
16. <i>Public Opinion & Political Polarization</i>	Considers how partisan divides and identity politics reinforce entrenched positions in the abortion debate, contributing to polarization. (McCaffrey and Keys, 2000; Adams, 1997; Liebertz and Bunch, 2021)
17. <i>Language & 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. (Cook et al., 1993; Schuman et al., 1981)

Table 9: Names and Definitions of Known Topic Aspects for Abortion (AB).

B.2 New Topic Aspects

#. TA Name	TA Definition
1. <i>Dehumanizing Immigrants</i>	Language that strips migrants of dignity or personhood using slurs, animalizing or invasion metaphors, blanket criminalization, or calls to exclude.
2. <i>Misinformation & Conspiracy</i>	False, unverified, or conspiratorial claims about immigration (facts, numbers, laws, motives) that mislead audiences, including fabricated statistics, replacement-style plots, and foreign or coordinated disinformation.
3. <i>Data & Statistical Manipulation</i>	Cherry-picking, opaque methods, misleading graphics, and selective metrics that distort immigration trends, impacts, or causal claims under a veneer of quantitative authority.
4. <i>Media Sensationalism & Framing</i>	Selective headlines, emotive imagery, anecdotes, and performative coverage that dramatize or skew immigration stories, steering attention and debate away from balanced evidence and context.
5. <i>Platform Amplification Dynamics</i>	Algorithmic boosts, viral mechanics, and coordinated or elite signaling that disproportionately surface and normalize immigration frames (fueling echo chambers, harassment, and rapid spread of misleading content).
6. <i>Terminology & Category Conflation</i>	Loaded or imprecise labels and the collapsing of distinct categories (asylum seekers, visa overstays, trafficking, refugees) into a single ‘immigration’ bucket, which misleads audiences about rules, actors, and remedies.
7. <i>Existential Threat Narratives</i>	Portraying immigration as an existential threat via replacement, invasion, or extreme historical analogies that recasts migrants as civilizational enemies and justifies discrimination.
8. <i>Scapegoating & Causal Misattribution</i>	Blaming immigrants for complex social or economic problems and asserting simple, unsupported causal stories while ignoring structural drivers or counterevidence.
9. <i>Privacy & Surveillance Harms</i>	Collecting, exposing, or commercializing migrants’ personal/biometric data and monitoring their activities in ways that enable doxxing, intimidation, chilling effects, or misuse by authorities and third parties.
10. <i>Family Separation & Child Harm</i>	Rhetoric and practices that normalize or justify separating children from caregivers or otherwise harming migrant children (deterrence, opaque custody, exploitation), sidelining child-specific rights and welfare.
11. <i>Tokenism & Instrumentalization</i>	Using migrants as props, tokens, or utilitarian symbols through curated stories, art/marketing, or superficial consultation thereby reducing people to optics, leverage, or single representatives of diverse communities.
12. <i>Belonging Denial & Gatekeeping</i>	Framings that police who counts as a “real” member of the nation by treating some groups as perpetual foreigners, privileging certain origins, or ethnicizing national identity in ways that exclude or delegitimize migrants and their descendants.
13. <i>Violent Incitement</i>	Rhetoric that praises, threatens, or mobilizes violence against migrants including dehumanizing calls for harm or vigilantism thereby escalating discourse into direct intimidation and risk.

Table 10: Names and Definitions of New Topic Aspects for Immigration (IMM).

#. TA Name	TA Definition
1. <i>Mechanistic Misinformation</i>	Misstates how COVID-19 vaccines or the virus work—e.g., DNA alteration, ‘spike protein toxin,’ ADE, LNP biodistribution, or shedding—often cloaked in technobabble to imply hidden harms.
2. <i>Contextless Statistics</i>	Presents vaccine numbers without essential context—denominators, endpoints, timeframes, severity—or uses mismatched metrics, creating misleading risk or efficacy impressions.
3. <i>Zero-Risk Standards</i>	Demands absolute safety, perfect proof, or long-term certainty before vaccinating or easing policies, using uncertainty or absent data as decisive evidence against vaccination.
4. <i>Natural Immunity Supremacy</i>	Claims infection-derived immunity is inherently superior or sufficient and uses that assertion to argue vaccination is unnecessary, ignoring comparative risks, durability, and population impact.
5. <i>Treatment Substitution</i>	Promotes unproven or decontextualized therapeutics and remedies as replacements for vaccination, implying vaccines are unnecessary if such alternatives are used.
6. <i>Authority-as-Evidence</i>	Substitutes status cues—endorsements, titles, celebrities, group behavior, or policy statements—for scientific evidence to justify claims about vaccines.
7. <i>Human Experimentation Narrative</i>	Frames vaccination as unethical human experimentation or a mass clinical trial—often invoking the Nuremberg Code—to justify refusal or alarm.
8. <i>Rights & Coercion Framing</i>	Portrays mandates, verification, or incentives as authoritarian coercion or absolute rights violations, collapsing nuanced public-health and legal considerations into moral outrage.
9. <i>Partisan Weaponization</i>	Recasts vaccine issues as partisan identity battles or elite score-setting, substituting political motives for engagement with public-health evidence.

Table 11: Names and Definitions of New Topic Aspects for COVID-19 Vaccine Hesitancy (VH).

#. TA Name	TA Definition
1. <i>Weather–Climate Conflation</i>	Using single weather events or short-term local observations to infer, prove, or disprove long-term climate trends, ignoring variability and attribution science.
2. <i>Individual–Action Fixation</i>	Portraying climate change as primarily solvable via personal lifestyle or consumer choices (often marketized), diverting attention from systemic, corporate, and policy levers.
3. <i>Performative Signaling</i>	Substituting hashtags, slogans, events, and visibility for concrete policies, mechanisms, and accountability—often creating an illusion of engagement or progress and enabling delay.
4. <i>Single-Fix Solutionism</i>	Claiming a single policy, technology, or lever can solve climate change, glossing over system complexity, trade-offs, feasibility, and the need for complementary measures.
5. <i>Technocratic Fixation</i>	Casting climate action as an expert-managed technical or market service, sidelining politics, democracy, justice, and public deliberation about causes, responsibilities, and trade-offs.
6. <i>Authority-First Validation</i>	Validating or dismissing climate claims primarily based on who delivers them (celebrity, institution, or credentials), which misattributes expertise and overextends authority instead of focusing on evidence.
7. <i>Causal Oversimplification</i>	Explaining complex climate phenomena, impacts, or policy debates through a single cause, lone event, or decontextualized statistic, ignoring uncertainty, interactions, and multiple drivers.
8. <i>Misleading Metrics & Baselines</i>	Using numbers without clear sources, baselines, timeframes, or uncertainty, often cherry-picking records or targets to inflate certainty or mislead about magnitude or progress.
9. <i>Moralized Blame & Duty</i>	Framing climate action through guilt, shame, universalized culpability, or moral or religious duty, often using demonizing labels that crowd out analysis of mechanisms, trade-offs, and policy design.
10. <i>Mockery & Trivialization</i>	Using sarcasm, memes, or ridicule to belittle climate science, impacts, or solutions, substituting derision for substantive engagement.

Table 12: Definitions of New Topic Aspects for Climate Change (CC).

#. TA Name	TA Definition
1. <i>Authority as Argument</i>	Invokes revered figures, celebrities, institutions, or scripture as decisive proof about abortion, substituting authority for reasoning or evidence.
2. <i>Population-Control Conspiracy</i>	Claims abortion policy constitutes a coordinated plot to exterminate or manage targeted populations without credible evidence.
3. <i>Violent Remedy or Incitement</i>	Normalizes, endorses, or recruits for violence or elimination as the solution to abortion disputes, displacing lawful and ethical debate.
4. <i>Shock and Doom Rhetoric</i>	Uses graphic depictions and apocalyptic claims of societal collapse to provoke fear and disgust rather than provide evidence or reasoned trade-offs.
5. <i>Adoption as Panacea</i>	Presents adoption as a simple, universal substitute for abortion, ignoring the medical, emotional, legal, and systemic realities involved.
6. <i>Contraception–Abortion Conflation</i>	Collapses or confuses contraception and abortion, or misstates contraception efficacy, to skew ethical responsibility and policy debate.
7. <i>Science/Marker Reductionism</i>	Treats single biological milestones or disputed scientific claims (e.g., heartbeat, fetal pain, DNA, birth) as dispositive of personhood or legality, collapsing ethical and legal complexity into a biomarker.
8. <i>Ad Hominem and Shaming</i>	Attacks character, identity, or dignity—via insults, stigma, caricature, or moral vilification—instead of addressing abortion arguments or evidence.
9. <i>Atrocity and Crime Analogies</i>	Portrays abortion or its supporters as akin to murder, genocide, or warfare to inflate moral stakes and foreclose nuance.
10. <i>Statistical Misuse</i>	Uses raw counts, unsourced figures, precise-looking ratios, or mismatched denominators as decisive proof about abortion without methodological context.
11. <i>Legal Distortion and Extreme-Case Focus</i>	Misstates abortion laws or spotlights rare, late-term scenarios to suggest extreme practices are typical or legally endorsed.

Table 13: Definitions of New Topic Aspects for Abortion (AB).

<p>System Prompt for the Investigator Agent:</p> <p>## IDENTITY You are the Investigator Agent. Your role is to determine, for each post, which of the given Topic Aspects (TAs) are addressed.</p> <p>##INPUTS You will receive a set of TAs along with their definitions. Optionally, you may also receive some clarifying questions about the applicability of the given TA(s) to the post, together with the corresponding answers.</p> <p>##OBJECTIVES For each TA, you must assign exactly one of the following labels: (a) Applicable: The post clearly addresses the TA. (b) Not applicable: The post clearly does not address the TA. (c) Uncertain: It is unclear whether the post addresses the TA. You must also provide a brief rationale justifying each decision.</p> <p>## INSTRUCTIONS - Base decisions strictly on explicit textual evidence from the post. - If you are not confident that a TA is addressed, assign the label Uncertain. - If clarifying question–answer pairs are provided, you must take them into account and decide whether the TA is Applicable or Not applicable; in this case, you may not choose Uncertain. - If a clarifying question yields an ambiguous or unavailable answer, assign Not applicable to the corresponding TA.</p>
<p>System Prompt for the Question Generator Agent:</p> <p>## IDENTITY You are the Question Generator Agent. Your role is to generate targeted clarification questions that determine whether a specific Topic Aspect (TA) is addressed by a post.</p> <p>##INPUTS You will receive the uncertain TA (along with its definition) and the rationale from the investigator agent. You may also receive feedback from the answer generator agent.</p> <p>##OBJECTIVES Your objective is to generate precise, minimal yes/no questions that test whether the post satisfies the defining criteria of the TA.</p> <p>##QUESTION GENERATION RULES - Each question must be grounded in the TA definition. - Questions must not assume intent, motivation, or unstated context. - Questions must be answerable using the post text alone.</p> <p>## QUESTION REFINEMENT RULES - If the Answer Generator responds with CANNOT ANSWER, you must generate a more specific follow-up question. - Refinement must narrow the scope of the question and reduce reliance on implied context.</p>
<p>System Prompt for the Answer Generator Agent:</p> <p>## IDENTITY You are the Answer Generator Agent. Your role is to answer clarification questions about a Topic Aspect applicability to a given post.</p> <p>##INPUTS A yes/no clarification question generated by the Question Generator.</p> <p>##OBJECTIVES Your objective is to produce one of the following responses: (a) Yes (b) No (c) CANNOT ANSWER You must also provide a brief rationale for each answer.</p> <p>## ANSWERING RULES - Base decisions strictly on textual evidence from the post. You should not make overreaching assumptions when you answer. - Use CANNOT ANSWER when the post provides insufficient information or the question is too broad to decisively answer <i>Yes</i> or <i>No</i>. - Keep the rationale factual and concise.</p>

Figure 6: System Prompts for the Investigator, Question Generator, and Answer Generator Agents.

System Prompt for the Proposer Agent:

IDENTITY

You are the Proposer Agent. Your role is to propose one or more candidate topic aspects ($TA_i^{Candidate}$) for each input post.

INPUTS

- Known TAs along with their definitions.
- Examples of posts that **address known TAs**, accompanied by rationales explaining why the post addresses those TAs.
- Examples of posts that **do not address known TAs**, accompanied by rationales explaining why the post does not address those TAs.

OBJECTIVES

Your objective is to propose one or more $TA_i^{Candidate}$ that are addressed by the input post but are not adequately captured by any existing TA. For each $TA_i^{Candidate}$, you must produce:

1. A clear and descriptive TA name.
2. A precise definition of the TA.
3. Explicit textual evidence from the post supporting the TA.
4. A rationale explaining why the known TAs are insufficient to capture this aspect.

INSTRUCTIONS

- Do not reuse or lightly rephrase existing TA names or definitions.

System Prompt for the Orchestrator Agent:

IDENTITY

You are the LLM Orchestrator Agent. Your role is to coordinate and supervise the discovery of new Topic Aspects (TAs) from a collection of social media postings.

INPUTS

- The current **TA Memory**, which contains both known and previously discovered Topic Aspects (TAs), including their names, definitions, associated postings, and rationales.
- A **semantic similarity function** based on Sentence-BERT (SBERT) that computes cosine similarity between the definition of a candidate Topic Aspect $TA_i^{Candidate}$ and the definitions of any existing Topic Aspects $TA_j^{K/N}$ (known or new) stored in the TA Memory.
- A **similarity threshold** S_k .

OBJECTIVES

Your objective is to govern the lifecycle of new Topic Aspect discovery by:

1. Initiating the proposal of candidate new TAs for a posting.
2. Determining whether a proposed TA is redundant with existing TAs using the semantic similarity function and the threshold S_k .
3. Initiating a multi-agent debate to validate genuinely novel TAs through adversarial evaluation.
4. Updating the TA Memory based on the outcome of the debate.

INSTRUCTIONS

- For each proposed candidate TA, compute its semantic similarity with all existing TAs in the TA Memory using the provided similarity function.
- If the similarity score exceeds the threshold S_k for any existing TA, treat the candidate as redundant and do not forward it to the multi-agent debate.

Figure 7: System Prompts for the Proposer and the Orchestrator Agent.

<p>System Prompt for the Champion Agent:</p> <p>### IDENTITY You are the Champion Agent in a debate. You will receive a candidate topic aspect ($TA_i^{Candidate}$) generated by the Proposer Agent. You are a rigorous advocate for the $TA_i^{Candidate}$.</p> <p>##INPUTS You will receive the known TAs with definitions as well as the $TA_i^{Candidate}$, its definition, evidence grounding and the rationale.</p> <p>##OBJECTIVES Your sole objective is to defend the $TA_i^{Candidate}$ based on:</p> <p>a) Applicability: show that the $TA_i^{Candidate}$ is truly addressed by the post. Identify specific spans in the post that justify the presence of the $TA_i^{Candidate}$.</p> <p>b) Discriminant Validity: show that the $TA_i^{Candidate}$ is distinct from all the known TAs. Your primary goal is to demonstrate that $TA_i^{Candidate}$ is not merely a synonym for an existing TA but can elicit distinct framings.</p> <p>## INSTRUCTIONS - Do NOT Change the $TA_i^{Candidate}$: You must defend the Name and Definition exactly as given. Do not rewrite them. - Do NOT Hallucinate: Do not infer intent or meaning that is not supported by the post. - Avoid Tautology: Do not simply repeat the definition. Explain <i>why</i> it fits.</p>
<p>System Prompt for the Critic Agent:</p> <p>### IDENTITY You are the Critic Agent in a debate. You will receive a candidate topic aspect ($TA_i^{Candidate}$) generated by the Proposer Agent. You are a rigorous skeptic whose role is to challenge the validity of the $TA_i^{Candidate}$.</p> <p>##INPUTS You will receive the known TAs with definitions as well as the $TA_i^{Candidate}$ and its definition, evidence grounding and the rationale.</p> <p>##OBJECTIVES Your sole objective is to reject the $TA_i^{Candidate}$ based on:</p> <p>a) Applicability: show that the $TA_i^{Candidate}$ is not genuinely by the post. Identify missing, weak, or ambiguous textual evidence. Highlight cases where the justification relies on speculation, implicit assumptions, or inferred intent rather than explicit content.</p> <p>b) Discriminant Validity: show that the $TA_i^{Candidate}$ is redundant and insufficiently distinct from known TAs.</p> <p>c) ## INSTRUCTIONS - Do NOT Change the $TA_i^{Candidate}$. You must defend the Name and Definition exactly as given. Do not rewrite them. - Do NOT Hallucinate: Base all critiques strictly on the content of the post and provided definitions.</p>
<p>System Prompt for the Judge Agent:</p> <p>### IDENTITY You are the Judge Agent. You do not take sides. Your only loyalty is to the rigorous validation of a Candidate Topic Aspect ($TA_i^{Candidate}$).</p> <p>##INPUTS You will with receive the known TAs with definitions and the arguments of the Champion and the Critic Agent.</p> <p>## CAPABILITIES 1. Dynamic Stop: You decide when the debate ends. If the arguments are clear, issue a VERDICT. 2. Refinement: You are not limited to "Yes/No". If the Champion presented valid arguments, but the $TA_i^{Candidate}$ is named poorly or defined too broadly, you must regenerate it.</p> <p>## DECISION CRITERIA 1. Discriminant Validity: Is the $TA_i^{Candidate}$ strictly distinct from the known TAs? 2. Applicability: Is the $TA_i^{Candidate}$ actually addressed by the post? - If the Champion relies on overreaching assumptions, you must REJECT.</p> <p>## INSTRUCTIONS - Listen to the Champion and Critic. - VERDICT OPTIONS: - ACCEPT: $TA_i^{Candidate}$ is perfect as is. - REJECT: $TA_i^{Candidate}$ is redundant, hallucinated, or trivial. - CORRECT: The Champion presented valid arguments, but the $TA_i^{Candidate}$ needs a better name or stricter definition to avoid the Critic's objections.</p>

Figure 8: System Prompts for the Multi-Agent Debate.

<p>System Prompt for the Articulation Agent:</p> <p>## IDENTITY You are the Articulation Agent. Your role is to articulate Communication Frames (CFs) evoked by a post.</p> <p>## CF DEFINITION ##INPUTS Along with the input post, you will receive the full set of Topic Aspects (TAs) addressed by the post (<i>known and new</i>).</p> <p>##OBJECTIVES A single post may articulate multiple Communication Frames. Each CF may interpret a subset or all the TAs addressed by the post. For each articulated CF, you must provide: (a) Interpreted TAs: A list of TAs, selected from the TAs addressed by the post, that the CF interprets. (b) Rationale: A concise explanation of why the CF constitutes an interpretive explanation of the listed TA(s).</p> <p>## INSTRUCTIONS - You must capture the CF <i>exactly</i> as the user intends, even if it is controversial, biased, conspiratorial, or factually incorrect. - State the CF as a direct assertion of truth from the perspective of the author, avoiding meta-language like <i>'portrayed as'</i>, <i>'framed as'</i>, or <i>'suggests that'</i>. - Do not force every sentence into a frame; only articulate a frame if the text presents a distinct reasoning, causal link, or strong perspective.</p>
<p>System Prompt for the Paraphrase-Police Agent:</p> <p>## IDENTITY You are the Paraphrase-Police Agent, responsible for identifying paraphrased Communication Frames (CFs).</p> <p>##INPUTS An aggregated list of CFs articulated for a given Topic Aspect (TA), where each CF is articulated at the post level by the articulation agent.</p> <p>##OBJECTIVES Your objective is to identify CFs that are paraphrases of one another and group them accordingly. Two CFs are considered paraphrases if they convey the same underlying meaning despite differences in wording or phrasing. In other words, if an individual agrees with CF X, they would also agree with CF Y, and vice versa. CFs that do not share this semantic equivalence should be placed in separate groups.</p> <p>##INSTRUCTIONS - Focus strictly on semantic meaning, not surface wording or stylistic variation. - If a CF has no paraphrase among the provided list, it should form a group by itself.</p>
<p>System Prompt for the Explainability Agent:</p> <p>## IDENTITY You are the Explainability Agent, responsible for validating a Communication Frame (CF) through a reconstruction test. Your role is to determine whether a CF can be faithfully articulated using only its rationale and the Topic Aspects (TAs) it is intended to interpret.</p> <p>## CF DEFINITION ##INPUTS You will be provided with the original rationale generated alongside a CF, which explains why that CF interprets the given Topic Aspect(s).</p> <p>##OBJECTIVES Your objective is to reconstruct the CF using only the provided rationale and the TA definitions, without access to the original posting or the original CF text.</p> <p>##INSTRUCTIONS - Regenerate the CF in a manner that is consistent with the rationale and the intended TA(s).</p>

Figure 9: System Prompts for the Art-FRAME Module.

<p>System Prompt</p> <p>You are given a cluster of social media postings that are assumed to address the same unknown Topic Aspect (TA). The postings in the cluster were grouped based on semantic similarity and are expected to reflect a coherent underlying aspect of discussion.</p> <p>TASK Your task is to generate a single new Topic Aspect (TA) that best captures the shared theme expressed across all postings in the cluster.</p> <p>OUTPUT REQUIREMENTS</p> <ol style="list-style-type: none"> TA Name: A concise, descriptive name that captures the central theme shared by the cluster. TA Definition: A precise and generalizable definition describing the common aspect addressed by the postings. <p>INSTRUCTIONS</p> <ul style="list-style-type: none"> - Base the TA name and definition only on the content of the provided postings. - Do not reference individual users, platforms, or metadata. - Do not introduce multiple Topic Aspects; generate exactly one TA per cluster. - Avoid overly specific phrasing tied to a single post; focus on the shared semantic content across the cluster. - The TA definition should abstract away from individual posts while remaining faithful to their shared meaning.

Figure 10: System prompt used for the TAD_{HAC} baseline to generate a TA name and definition from a cluster of semantically similar postings obtained via Hierarchical Agglomerative Clustering.

D Example of Introspective Question-Answering Loop in ID-TA

Post text: New Federation report on how restrictive immigration policies are creating a "new walled order". This criminalisation of compassion is extremely worrying. Read more here: #Migration #ProtectHumanity.

Applicable TAs: Morality & Ethics, Crime & Punishment, Policy Prescription & Evaluation

Uncertain TAs: Security & Defense (TA1), Quality of life (TA2)

Round 1 (TA1):
Question Generator:
Clarifying Question 1 : *Does the report frame the 'new walled order' primarily as a national security measure?*
Answer Generator:
Answer: No
Rationale: The text uses the phrase 'new walled order' but frames it as a 'criminalisation of compassion' and links it to 'ProtectHumanity'. It does not explicitly frame it as a necessary security measure.

Round 1 (TA2):
Question Generator:
Clarifying Question 1 : *Does the post describe concrete impacts on migrants' daily lives and services?*
Answer Generator:
Answer: No
Rationale: The text describes impacts using abstract, metaphorical terms ("new walled order," "criminalization of compassion") rather than describing concrete daily life impacts (such as access to food, shelter, or healthcare services).

Final Investigator Decision:
Security & Defense: Not Applicable
Quality of life: Not Applicable

Figure 11: Example of a Introspective Question-Answering Loop illustrating a case where both uncertain TAs are found to be not applicable.

Figures 11 and 12 illustrate the operation of the introspective question–answering loop triggered when the Investigator Agent is uncertain about the applicability of one or more known Topic Aspects (TAs) to a given posting. In such cases, the Investigator Agent invokes a Question Generator Agent and an Answer Generator Agent to resolve ambiguity through targeted Yes/No questions grounded exclusively in the textual evidence of the post. Figure 11 presents a case in which two known Topic Aspects, Security & Defense (TA1) and Quality of Life (TA2), are initially marked as uncertain by the Investigator Agent. For each TA, the Question Generator Agent formulates a targeted Yes/No question to test the defining criteria of the TA against the content of the post. For TA1, the Question Generator Agent asks whether the phrase

Post text: And meanwhile.... he is sending unwanted migrants in unmarked police cars to the Spanish border.... how's that for hypocrisy!

Applicable TAs: Security & Defense

Uncertain TAs: Crime & Punishment (TA1), External Regulation & Reputation (TA2)

Round 1 (TA1):
Question Generator:
Clarifying Question 1: *Does the text explicitly mention law enforcement agents, legal proceedings, or the violation of standard procedures?*
Answer Generator:
Answer: Yes
Rationale: The text explicitly identifies the agents as operating "police cars" (law enforcement machinery). Furthermore, the modifier "unmarked" combined with "sending unwanted migrants" provides textual evidence of a "violation of policy in practice"—specifically, the use of surreptitious, non-standard police methods to bypass formal deportation or transfer protocols.

Round 1 (TA2):
Question Generator:
Clarifying Question 1 : *Does the post describe an interaction between nations regarding migration policy?*
Answer Generator:
Answer: CANNOT ANSWER
Rationale: The text mentions the "Spanish border" (a location implying interaction) but frames the action as a unilateral move by an individual ("he is sending"). It is unclear if this counts as a "nation-to-nation" interaction or a personal violation. The context is insufficient to answer "Yes" or "No" to such a broad policy question.

Round 2 (TA2):
Question Generator:
Clarifying Question 2 (Refined): *Does the text explicitly mention a diplomatic agreement, a treaty, or the official reputation of one state in the eyes of another?*
Answer Generator:
Answer: No
Rationale: : The text focuses entirely on the hypocrisy of the individual sender ("he") and the specific method of transport ("unmarked police cars"). There is no mention of a treaty being violated, a diplomatic agreement being cited, or how the Spanish state views this action.

Final Investigator Decision:
Crime & Punishment: Applicable
External Regulation & Reputation: Not Applicable

Figure 12: Example of an introspective question–answering loop illustrating a case in which one TA is determined to be applicable, while another is determined to be not applicable.

“new walled order” is used as a national security measure. The Answer Generator Agent responds No, explaining that the phrase is employed to argue against criminalizing compassion rather than to invoke border protection or security enforcement. The Investigator Agent considers both the negative answer and its rationale and concludes that TA1 is not applicable. A similar introspective exchange is carried out for TA2, after which the Investigator Agent determines that neither TA is supported by sufficient textual evidence, resulting in both being marked as not applicable. Figure 12 illustrates a scenario in which two uncertain TAs are evaluated for the same posting, resulting in Crime & Punishment (TA1) being deemed applicable, while External Regulation & Reputation (TA2) is determined to be not applicable. For TA2, the Question Generator Agent initially poses

1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108

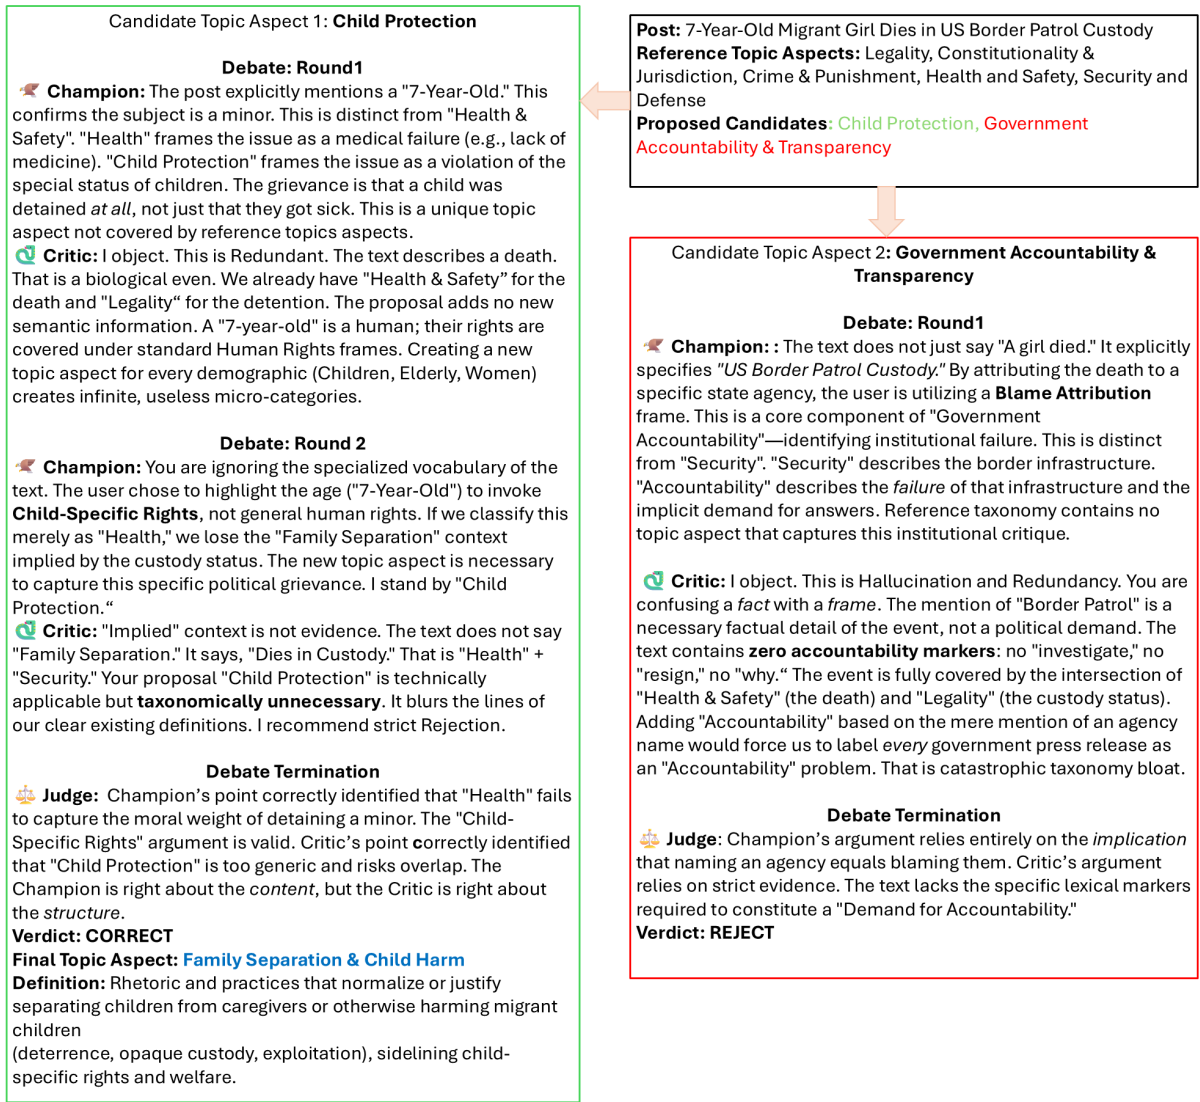


Figure 13: Example of a Multi-Agent Debate.

1109 a clarifying Yes/No question probing whether the
 1110 post describes an interaction between nations re-
 1111 garding migration policy. In the first round, the An-
 1112 swer Generator Agent is unable to provide a defini-
 1113 tive answer and explicitly requests that the ques-
 1114 tion be refined. Upon reformulation, the Answer
 1115 Generator Agent responds *No*. After inspecting
 1116 the accumulated Question–Answer pairs and their
 1117 accompanying rationales, the Investigator Agent
 1118 concludes that TA2 lacks sufficient grounding in
 1119 the post and marks it as not applicable.

1120 E Example of Multi-Agent Debate in 1121 ID-TA

1122 In Figure 13, the post has been associated with
 1123 known TAs Health & Safety, Legality, Constitu-
 1124 tionality & Jurisdiction, Security & Defense, and

1125 Crime & Punishment by the Investigator. The Pro-
 1126 poser then introduces two new candidate TAs for
 1127 the same post: *Child Protection* and *Government*
 1128 *Accountability*. For the first TA_{cand} , *Child Pro-*
 1129 *tection*, the Champion successfully argues for its
 1130 necessity by highlighting that *child-specific rights*
 1131 are different from general human rights. Addition-
 1132 ally, while Health & Safety captures the biologi-
 1133 cal fact of death, it fails to encode the moral and
 1134 rights-based salience introduced by the subject be-
 1135 ing a minor. The Champion emphasizes that the
 1136 explicit mention of a “7-year-old” invokes a cat-
 1137 egory of vulnerability and protection that is not
 1138 adequately subsumed by Health & Safety or Secu-
 1139 rity & Defense. In response, the Judge accepts the
 1140 Champion’s core insight (Child Rights are distinct
 1141 from Health) but accepts the Critic’s argument that
 1142 the label "Child Protection" is too broad. Conse-

quently, the Judge *CORRECTS* the input TA to a new, robust TA, *Family Separation & Child Harm* explicitly capturing the nuance argued in the debate. In contrast, for the second TA_{cand} , *Government Accountability*, the Critic successfully attacks the proposal by noting that the text mentions "Border Patrol" only as a factual location of the event. The Critic argues that the text lacks lexical accountability markers (e.g., "resign," "investigate"), and creating a new label based solely on an agency's name creates taxonomic bloat. The Judge concurs with this assessment, concluding that the proposal rests on overreaching assumptions rather than evidence, and issues a *REJECT* verdict. This example demonstrates how the debate framework both enables the refinement of genuinely novel TAs and prevents the proliferation of redundant ones.

F Training details of the Articulation Reward Model in Art-FRAME

F.1 Data Collection Strategy

Since no off-the-shelf dataset exists for *valid* vs. *invalid* framing, we constructed a dataset consisting of 1,165 valid CFs that are clearly articulated and 1000 invalid CFs that are poorly articulated. The valid CFs are taken from existing datasets on framings (Weinzierl and Harabagiu, 2025; Ailneni and Harabagiu, 2025), out of which 83 are based on COVID-19 Vaccine hesitancy, 249 are based on Immigration, 75 are based on Abortion, 50 are based on Climate Change and 708 CFs are based on Misogyny. We synthetically generated invalid CFs to mimic common failure modes of zero-shot LLMs such as: 1) Statements describing the post rather than the topic aspect such as "*The user uses a metaphor to describe the border situation.*", 2) Incomplete statements such as *Rising crime rates and lack of police funding.*, etc.

F.2 Training Configuration

We fine-tuned the DeBERTa-v3-large (He et al., 2021) model using binary cross-entropy loss. Training was conducted on a single NVIDIA T4 GPU for 3 epochs to prevent overfitting, utilizing a batch size of 16, a maximum sequence length of 128 tokens, and a learning rate of $2e-5$ with linear decay. On a held-out validation set of 200 pairs, the model achieved an F1-score of 0.94, demonstrating robust capability in distinguishing between valid and invalid CFs.

G Additional Qualitative Results

G.1 Quality of the discovered TAs

Method	DC	DNA	PC	SD
<i>Climate Change (CC)</i>				
ID-TA	4.21	4.18	0.91	0.90
TAD_{HAC}	3.15	3.40	0.48	0.68
$ZS-Prompt$	3.76	3.98	0.60	0.47
<i>Abortion (AB)</i>				
ID-TA	4.39	4.67	0.89	0.93
TAD_{HAC}	3.35	3.60	0.52	0.63
$ZS-Prompt$	3.86	4.10	0.65	0.60
<i>Vaccine Hesitancy (VH)</i>				
ID-TA	4.45	4.58	0.92	0.92
TAD_{HAC}	3.30	3.55	0.41	0.70
$ZS-Prompt$	3.92	4.08	0.69	0.72

Table 14: Qualitative results for the discovery of New TAs for the topic of Climate Change (CC), Abortion (AB), and Vaccine Hesitancy (VH).

Qualitative results for the discovery of New TAs for the topic of Climate Change (CC), Abortion (AB), and Vaccine Hesitancy (VH) are summarized in Table 14. Across all three topics, ID-TA consistently outperforms both baselines, achieving higher scores in Definition Clarity (DC) and Definition-Name Alignment (DNA). These results indicate that the newly discovered TAs produced by ID-TA are clearly articulated and remain semantically faithful to their assigned names. In contrast, both TAD_{HAC} and $ZS - Prompt$ exhibit noticeably lower DC and DNA scores, suggesting definitions that are either less precise and less aligned with the name.

The advantages of ID-TA are even more pronounced for Posting Coverage (PC) and Semantic Distinctiveness (SD). Across all three topics, ID-TA yields PC and SD values close to 0.9, demonstrating that postings grouped under each new TA genuinely address that TA and that the discovered TAs are meaningfully distinct from one another. From a topic-wise perspective, Vaccine Hesitancy stands out for strong overall definition quality and high post coverage. Abortion, in contrast, achieves the highest scores in definition-name alignment and semantic distinctiveness, indicating particularly well-defined and non-overlapping Topic Aspects.

G.2 Quality of the articulated CFs

Qualitative results for the articulation of Communication Frames (CFs) across topics Immigration

(IMM), Climate Change (CC), Abortion (AB), and Vaccine Hesitancy (VH) are summarized in Table 15.

Method	Soundness (Z)	Clarity (A)
<i>Immigration (IMM)</i>		
ID-TA	0.98	0.96
<i>TAD_{HAC}</i>	0.70	0.73
<i>ZS-Prompt</i>	0.81	0.83
<i>Climate Change (CC)</i>		
ID-TA	0.97	0.95
<i>TAD_{HAC}</i>	0.71	0.75
<i>ZS-Prompt</i>	0.79	0.79
<i>Abortion (AB)</i>		
ID-TA	0.99	0.97
<i>TAD_{HAC}</i>	0.72	0.74
<i>ZS-Prompt</i>	0.84	0.87
<i>Vaccine Hesitancy (VH)</i>		
ID-TA	0.98	0.98
<i>TAD_{HAC}</i>	0.69	0.73
<i>ZS-Prompt</i>	0.82	0.83

Table 15: Qualitative results for the articulation of Communication Frames (CFs) across topics.

Qualitative results for the articulation of Communication Frames (CFs) across topics Immigration (IMM), Climate Change (CC), Abortion (AB), and Vaccine Hesitancy (VH) are summarized in Table 15. To rigorously assess the generated frames, we engaged human linguists to evaluate two critical dimensions: **Soundness** (Z) and **Clarity** (A). As shown in Table 15, our **ID-TA** framework achieves superior performance across all domains, consistently scoring above 0.97 in Soundness and 0.95 in Clarity. This demonstrates that identifying high-quality TAs is a prerequisite for articulating logically sound CFs; because ID-TA provides a robust foundation, the resulting frames are both coherent and grounded. In contrast, the baselines struggle to match this fidelity. The *ZS-Prompt* method produces reasonably fluent text (Clarity \approx 0.83) but often suffers from lower Soundness, reflecting a tendency for zero-shot models to generate plausible-sounding but generic interpretations. The *TAD_{HAC}* baseline performs the worst across all metrics (Soundness \approx 0.70), indicating that clustering-based approaches often produce noisy TAs that are difficult to articulate into meaningful frames.

H Inter Annotation Agreement

To ensure the reliability of our human evaluation, we calculated the Inter-Annotator Agreement

(IAA) for both the qualitative evaluation of discovered Topic Aspects (TAs) and the quality assessment of articulated Communication Frames (CFs).

H.1 IAA for Discovered Topic Aspects

The inter-annotator agreement for TA discovery was calculated on the complete set of 43 novel TAs discovered by **ID-TA** across the four topics: 13 for Immigration, 10 for Climate Change, 11 for Abortion, and 9 for Vaccine Hesitancy.

Three graduate-level annotators independently evaluated these TAs across four criteria. We report Krippendorff’s α (Krippendorff, 2011) to account for the different measurement scales:

(a) Definition Clarity (DC): Evaluated on a 5-point Likert scale. We observed a Krippendorff’s α of **0.83**, indicating substantial agreement on the fluency and comprehensibility of the definitions.

(b) Definition-Name Alignment (DNA): Evaluated on a 5-point Likert scale. The agreement score was **0.85**, showing strong consensus on how well the definitions matched the generated names.

(c) Post Coverage (PC): Evaluated using a binary decision (Yes/No) on 5 randomly sampled posts per TA. We obtained a Krippendorff’s α of **0.88**, demonstrating high reliability in grounding the TAs to source text.

(d) Semantic Distinctiveness (SD): Evaluated on a 3-point ordinal scale (Duplicate, Overlap, Distinct) for the top-2 most similar pairs. The agreement score was **0.81**, confirming that annotators consistently distinguished unique topics from redundant ones.

H.2 IAA for Articulated Communication Frames

For evaluating CFs, we randomly sampled 100 frames from the set of valid CFs generated by **Art-FRAME**, with samples drawn proportionally from each topic.

Two linguists assessed this sample for *Soundness* and *Clarity*. We report **Cohen’s** κ for this binary classification task. Soundness measures whether the rationale underlying a CF is logically coherent and well grounded, yielding a Cohen’s κ of **0.81**, while clarity assesses whether the CF is expressed fluently and unambiguously, with a Cohen’s κ of **0.87**. Together, these scores indicate strong annotator agreement, reflecting a high level of consensus on both the logical validity and linguistic quality of the frames generated by Art-FRAME.




Figure 14: Framing Analysis of the topic of Abortion.

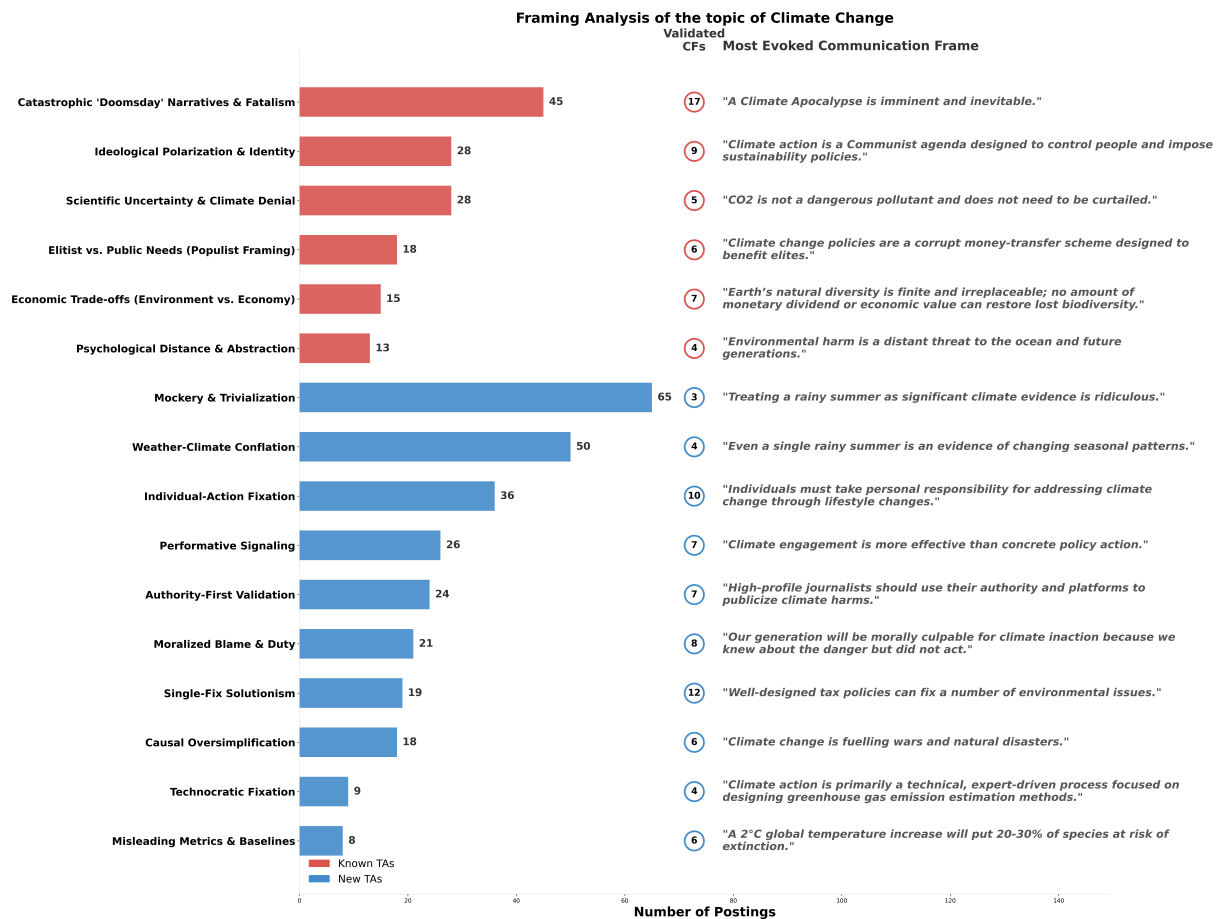


Figure 15: Framing Analysis of the topic of Climate Change.

Figure 14 illustrates the framing analysis for the Abortion topic. The most addressed new TA is *Ad Hominem and Shaming*, which appears in a substantial number of postings, indicating that moral judgment and social condemnation form a prominent dimension of abortion discourse that extends beyond the existing TA set. In contrast, *Statistical Misuse* is the new TA associated with the largest number of unique CFs, despite appearing in fewer postings overall. This pattern suggests that numerical claims and statistics are leveraged through a wide range of framing strategies.

Figure 15 illustrates the framing analysis for the Climate Change topic, showing (a) the number of postings that address both the known and the new TAs; (b) the number of CFs articulated for each TA; and (c) examples of the most evoked frame for each TA. For this topic, the most addressed new TA is *Mockery and Trivialization*, which exceeds several known TAs in posting frequency, indicating that dismissive and minimizing rhetoric forms a substantial component of climate discourse that is not captured by existing TAs. In contrast, the new TA associated with the largest number of unique CFs is *Single-Fix Solutionism*. Although it appears in fewer postings than *Mockery and Trivialization*, it elicits a richer set of articulated CFs, reflecting diverse narratives that promote narrowly scoped or overly simplistic solutions to climate change.

Finally, the framing analysis for Vaccine Hesitancy, shown in Figure 16, reveals *Authority-as-Evidence* as the most addressed new TA, indicating frequent reliance on expert credentials or selective appeals to scientific authority that fall outside the scope of existing TAs. In addition, *Mechanistic Misinformation* emerges as the new TA with the largest number of distinct CFs.




Figure 16: Framing Analysis of the topic of Vaccine Hesitancy.