

Retain or Reframe? A Computational Framework for the Analysis of Framing in News Articles and Reader Comments

Matteo Guida Yulia Otmakhova Eduard Hovy Lea Frermann

School of Computing and Information Systems,

The University of Melbourne

guida@student.unimelb.edu.au

{y.otmakhova, eduard.hovy, lea.frermann}@unimelb.edu.au

Abstract

When a news article describes immigration as an “economic burden” or a “humanitarian crisis,” it selectively emphasizes certain aspects of the issue. Although this *framing* shapes how the public interprets such issues, audiences do not absorb frames passively but actively reorganize the presented information. While this relationship between source content and audience response is well-documented in the social sciences, NLP approaches often ignore it, analyzing frames in articles and responses in isolation. We present the first computational framework for large-scale analysis of framing across source content (news articles) and audience responses (reader comments). Methodologically, we refine frame labels and develop a framework that reconstructs primary frames in articles and comments from sentence-level predictions, and aligns articles with topically relevant comments. Applying our framework across eleven topics and two news outlets, we find that frame reuse in comments correlates highly across outlets, and that readers often selectively engage with frames of the articles. We release a frame classifier that performs well on both articles and comments, a dataset of article and comment sentences manually labeled for frames, and a large-scale dataset of articles and comments with predicted frame labels.¹

1 Introduction

Framing research has long recognized that the way news content is presented fundamentally shapes public understanding of issues (Entman, 1993; Chong and Druckman, 2007). However, audiences do not passively absorb frames but rather “actively

filter, sort and reorganize information in personally meaningful ways in the process of constructing an understanding of public issues” (Neuman et al., 1992, pg.77). The constructionist understanding of framing emphasizes that frames cannot be fully understood without considering the interaction between source and receiver (Scheufele, 1999; Gorp, 2007; Hubner and Dixon, 2023), as part of an interactive and dynamic process in which frames activate and resonate (or fail to) with personal schemas through active meaning construction and interpretation. This is particularly evident on news websites with comment threads that create immediate opportunities for readers to publicly respond to framed content – a paradigm shift from one-way communication to participatory discourse where audiences actively engage with and reconstruct news content (Risch and Krestel, 2020).

However, current research has yet to fully explore the relationship between media framing and audience reception. Cognitive studies have demonstrated that the framing of news articles can significantly shape readers’ perceptions of an issue in controlled, small-scale experiments (Price et al., 1997; Valkenburg et al., 1999; Nelson and Oxley, 1999). Communication scholars have examined the influence of framing by analyzing online news articles and their associated reader comments (Zhou and Moy, 2007; Muñiz et al., 2015; Dargay, 2016; Hubner and Dixon, 2023). These studies often focused on a single issue and relied on small, manually annotated data sets, which limited generalizability and led to inconsistent findings.

Meanwhile, Natural Language Processing (NLP) approaches have developed numerous classification systems (for a review, see Ali and Hassan (2022)) aimed at detecting frames in text of different lengths such as news articles (Field et al., 2018), or social media posts (Mendelsohn et al., 2021). Yet these large-scale methods treat frame detection in source documents (such as news) and

¹We released our data at: <https://github.com/mattguida/FrAC>; and fine-tuned frame classifier: <https://huggingface.co/mattdr/sentence-frame-classifier>




Figure 1: Overview of Our Framework. Bottom: Article with two comments illustrating **reframing** (Comment 1; changes the article frame) and **frame retention** (Comment 2; keeps the article frame).

responses to them (such as comments or social media posts) as isolated classification tasks without connecting the framed source text and the framing in the audience response (Otmakhova et al., 2024).

If framing is truly a process in which frames are actively constructed and interpreted by an audience (Scheufele, 1999; Gorp, 2007; Hubner and Dixon, 2023), then the audience’s response is not merely a posited outcome, but a fundamental part of the analysis that has been largely overlooked or observed at a limited scale. To address this gap, we propose a computational framework that investigates at scale the relationship between how information is framed in news articles and how readers respond to such framed content in associated comments. Using an established framing taxonomy (Boydston et al., 2014; Card et al., 2015), we examine both source content and audience responses as interconnected elements of a broader communicative process (Hubner and Dixon, 2023).

Throughout this paper, we use the phrase **media framing effects** to describe the relation between news article framing and that of reader comments, as an umbrella term that covers two phenomena. **Frame retention** occurs when readers adopt the primary frame² used in the source article in their comments. **Reframing** occurs when readers shift to a different frame in their comments: either a secondary frame in the article (*selective* reframing), or a completely new, unmentioned frame (*complete* reframing). Figure 1 illustrates this with an example.

Our large-scale analysis across eleven diverse topics and two news outlets (The New York Times and The Globe and Mail) reveals systematic patterns of media framing effects. We find that (a)

²We define the primary frame as the one that covers the most sentences of an article, but at least 40% (see Section 6.2).

frame retention rates vary by source document frame, and this variation correlates highly across the two outlets, with value-laden frames being more often reframed; (b) retention patterns vary systematically by topic, and (c) audiences engage in selective reframing, adopting secondary frames rather than completely accepting or rejecting primary frames.

We make three main contributions: (1) a computational framework for investigating the relationship between source frames (articles) and frames in associated responses (comments), (2) a frame classification model that effectively generalizes across different text types (news articles and reader comments) and topics, and (3) two data sets: a manually labeled evaluation corpus (FrAC) and a large-scale data set of articles and comments aligned by topic group and labeled with primary frames, for reproducibility and supporting future research.

2 Related Work

Cognitive Studies on Framing Effects Cognitive research has established that news frames significantly influence audience interpretation using controlled experimental designs, where participants were assigned to different framing conditions of the same issue and asked to report their thoughts. Valkenburg et al. (1999) examined how news stories framed as conflict, human interest, attribution of responsibility, or economic consequences affected readers, finding that participants generally adopted news frames. However, Price et al. (1997) showed that frames not only guide interpretation but also stimulate new thoughts and feelings not explicitly present in the original content. Nelson and Oxley (1999) and Shen (2004) showed that framing primarily influences the importance of existing beliefs rather than changing them, with frames being the most effective when they resonate with pre-existing schemas. These studies were restricted to small-scale experimental settings and narrow demographic representation, and thus could not generalize to natural media environments with direct audience responses.

Media Framing Effects in Communication Studies Communication researchers have attempted to bridge this gap by analyzing reader comments on online news articles, albeit with mixed results. Several studies have found evidence of frame retention. Kleut and Milojevic (2021) analyzed online coverage of Serbian protests and showed that comments

were influenced by the article’s frames. Dargay (2016) examined terrorism coverage in The New York Times, finding that commentators engaged with primary frames of the articles, although they also introduced new frames.

Other studies have found no evidence of retention. Holton et al. (2014) analyzed online health coverage across major outlets and found no evidence of frame retention in comment threads. Muñiz et al. (2015) examined coverage of an oil company expropriation in Spanish online outlets and found that media frames did not produce corresponding effects in reader comments. Similarly, Zhou and Moy (2007) found no framing effects in news about Chinese sociopolitical incidents. Finally, in climate change coverage, Koteyko et al. (2013) found minimal retention of frames in reader comments on online articles from UK outlets, while Hubner and Dixon (2023) showed that retweets sharing climate news rarely replicated the original frames.

There are several explanations for these inconsistencies, including small data sets, focus on a single topic, and the different adopted framing conceptualizations. In this study, we increase the data set size, keep the framing conceptualization fixed but apply it across eleven topics, in order to disentangle some of these confounds.

Computational Approaches NLP approaches to framing have evolved from exploratory topic-modeling (DiMaggio et al., 2013) to classification tasks (for a review, see Ali and Hassan (2022)). The Media Frame Corpus (MFC, Card et al. (2015)), based on Boydstun et al. (2014), has been particularly influential, spurring various frame classification methods including probabilistic approaches (e.g., Johnson and Goldwasser (2018)), neural networks (e.g., Naderi and Hirst (2017)), and pre-trained language models (e.g., Khanehazar et al. (2021)). These computational approaches have analyzed frames at different granularities: headline level (Liu et al., 2019a; Akyürek et al., 2020), document level (Field et al., 2018; Ji and Smith, 2017; Piskorski et al., 2023), span level (Khanehazar et al., 2021), sentence level (Naderi and Hirst, 2017; Hartmann et al., 2019), and paragraph level (Roy and Goldwasser, 2020). Some studies have investigated the generalizability of the MFC to other domains and languages: Khanehazar et al. (2019) found reduced performance when applying fine-tuned BERT-based classifiers to different En-

glish contexts (Australian parliamentary speeches), while Daffara et al. (2025) tested cross-linguistic transfer using both fine-tuned pre-trained language models and zero-shot LLM classification on Portuguese news articles.

However, despite acknowledging framing’s theoretical impact on audiences, these existing computational approaches treat frame detection in source content and audience responses as isolated classification tasks, ignoring the interactive process that defines constructionist framing theory.

3 Research Questions

Our goal is to characterize and quantify the *constructivist view on framing* at scale. We formulate research questions that have been raised in the social sciences, albeit on a small scale often regarding a single topic and a specific framing inventory which makes it difficult to generalize results across scenarios. With orders of magnitude more data that covers a range of topics and is analyzed with a unified framing inventory, we are well positioned to answer the following research questions.

RQ1: To what extent are primary article frames retained in reader comments?

We first focus on the extent to which the most prominent frame in the article is reflected in the reader comments. From a constructivist point of view, this assesses the extent to which a chosen article frame *sticks with* the readership.

Next, we address the fact that prior work was often topic-specific and led to inconclusive results about the extent of frame retention. Here, we can quantify the effect of topics on frame retention within a unified framework. We ask:

RQ2: How does frame retention vary across topics?

Finally, we examine *how* readers reconstruct media messages when they do reframe (i.e., they choose not to retain the primary article frame). Our framework allows us to paint a more nuanced picture on reframing than previous work. First, in line with prior studies we quantify how often readers adopt a frame that was never used in the article (*complete reframing*). In addition, we also explore how often readers shift from the primary article frame to a secondary article frame (*selective reframing*). This leads to our final research question:

RQ3: If readers reject the primary article frame, how do they reframe?

We next introduce our framework and methodology, before we address our research questions in Section 7.

4 Framework

To address these research questions, we propose a scalable and automated framework for analyzing media framing effects (retention and reframing).

4.1 Data

Our framework requires a data set of source documents d_i (e.g., news articles, headlines, tweets, etc.) each paired with a set of response documents $R_{d_i} = \{r_1, r_2, \dots\}$ (e.g., comments, retweets, blog posts, etc.) directly linked to the source ($D = \{(d_i, R_{d_i})\}_{i=1}^N$). For the purpose of this study, we use two existing data sets of news articles with their associated comments:

- **SOCC** (The SFU Opinion Comments Corpus; Kolhatkar et al. (2020))³: consists of 10,339 articles and over 1.2 million comments sourced from The Globe and Mail, a major Canadian newspaper. The corpus covers five years (2012–2016).
- **NYT** (New York Times): consists of articles and comments from *The New York Times*⁴ with over two million comments associated with approximately 9,000 articles published between January-May 2017 and January-April 2018.

This pairing of source documents with response documents directly operationalizes the constructionist understanding of framing. In our analysis, news articles represent the “source framing” – how journalists and media outlets choose to present issues through particular frames. Reader comments represent the “audience response” – how readers actively interpret, accept, modify, or reject those frames based on their own cognitive schemas and values. We focus exclusively on top-level comments that directly respond to articles, excluding nested replies within comment threads, as we want to capture direct audience responses to the original framed content.

³Available at: <https://www.kaggle.com/datasets/mtaboda/sfu-opinion-and-comments-corpus-socc>

⁴Available at: <https://www.kaggle.com/datasets/aashita/nyt-comments>. Article texts were retrieved using the article titles in the data set via ProQuest.

4.2 Framework Structure

The framework consists of three core components that work in concert to analyze framing dynamics between articles and comments (Figure 1):

1. **Frame Classification:** To compare framing across source documents and responses, we build on established framing taxonomies to develop and evaluate frame classifiers that work reliably across both articles and comments.
2. **Primary Frame Detection:** In line with prior work in cognitive and communication studies (Section 2), we identify and compare a single, *primary* frame (defined as the most recurring one) in source documents and responses. We thus require a method to determine the primary frame in each input document. We define the primary frame as the most frequently occurring sentence-level frame within a document (see Section 6.2 for details). In doing so, we can distinguish the retention of the primary article frame from *selective reframing* into a *secondary frame* of the original article.
3. **Topic Identification:** We assign each source document and reader response to a topic, for two reasons. First, assigning topics to articles allows us to perform analysis of framing effects over multiple topics, advancing over previous, topic-specific studies (Section 2) and potentially explaining their divergent results. Second, we want to capture framing effects only for responses where commentators stay on topic with the source, rather than digress into a different topic (with potentially a different frame, which can lead to underestimation of retention rates).

Below, we first introduce a new corpus that we annotated for the *frame classification* subtask (Section 5). Section 6 explains and evaluates our implementations of the three framework components introduced above. Finally, we apply the framework to construct a data set (Section 6.4) that we use to analyze media framing effects (Section 7).

5 The FrAC Corpus

To apply our framework, we require (1) a frame label set that works robustly across both source documents (news articles) and reader reactions (comments), and across a variety of different topics; and (2) a way to validate predicted frames across these two types of documents. To this end, we introduce the **Frames in Articles and Comments (FrAC)** cor-

pus. Our corpus uses a more robust frame label set derived from that of the Media Frames Corpus (Section 5.1), and has high-quality manual annotations showing that humans can reliably apply our label set to both document types (Section 5.2).

5.1 Frame Labels and Granularity

We devise a robust frame label set, building on the Media Frames Corpus (MFC), a foundational data set in computational framing analysis based on [Boydston et al. \(2014\)](#)’s theoretical framework operationalizing framing through 15 generic frame categories. The original MFC corpus ([Card et al., 2015](#)) consists only of news articles; however, more recently the framework has been successfully applied to other types of documents, including online forum comments ([Hartmann et al., 2019](#)), suggesting suitability for our task that involves both long, formal documents and short, less formal responses.

First, to apply the MFC labels across articles and comments, we need to define a granularity that is meaningful for both of them. While most NLP studies on the MFC modeled article-level primary frames, this is not feasible for reader responses which are often brief and composed of single sentences. We therefore chose the sentence level – the largest shared textual unit between the two types of text – as our basic unit of analysis.⁵

Second, we improve the robustness of the MFC label set as it has been found (e.g., [Ali and Hasan \(2022\)](#); [Daffara et al. \(2025\)](#)) that some MFC frames are poorly defined and difficult to distinguish, leading to low inter-annotator agreement and model performance.

To do so, we isolated highly confused frame labels in the original MFC data set by first identifying span-level MFC annotations (of ≥ 3 words) to which two or more annotators assigned different frame labels. We examined such inconsistent annotations using confusion matrices – both for all MFC data and separately for each topic included in the MFC corpus. We took a conservative approach by only merging frame labels that were confused consistently across all topics,⁶ resulting

⁵Section 6.2 explains how we derive a primary document frame from sentence-level frames.

⁶For example, “Health and Safety” and “Legality, constitutionality and jurisprudence” labels are frequently assigned to the same spans in “Tobacco” and “Gun Control”, but not in other topics. In this case, we conclude that the confusion is due to the topic specifics rather than meaning of the label, and keep such semantically distinct labels separate.

1) Legality and Crime*, 2) Political and Policies ⁺ , 3) Economic, 4) Health and Safety, 5) Cultural Identity, 6) Public Opinion, 7) Morality, 8) Fairness and Equality, 9) Security and Defense, 10) Other.

Table 1: The FrAC Frame Label Set. Label mergers: *= {“Legality, constitutionality and jurisprudence”, “Crime and punishment”}; += {“Policy prescription and evaluation”, “Political”}.

in two merges of two frames (see Table 1 and Table 10 in the Appendix for a list of the frames and definitions). To further improve robustness of the label set, we removed the two least frequent classes (“External regulation and reputation” and “Capacity and resources”) which accounted for less than 0.7% of the data. We finally obtained 9 frame categories, plus a generic “Other” frame to account for instances that do not match any of our nine frames. This category captures both genuinely unframed content and framed content that falls outside our nine established categories.

5.2 Human Annotation

To validate our frame label set (and model predictions) we collected human annotation on a subset of the SOCC corpus (Section 4.1).

Annotation Setup A custom web-based annotation interface was developed using FastAPI.⁷ The interface included detailed instructions and examples, and an interactive training phase with immediate feedback and explanations. See Appendix E for detailed descriptions.

We recruited six native English speakers as annotators, all with a background in NLP or political science and journalism.⁸ Each annotator was presented with 100 sentences from the SOCC corpus, drawn evenly from articles and comments, 20% of the sentences overlapped with another annotator to assess agreement. One additional annotator (an author of this paper) annotated all sentences, so all sentences in the resulting data set had at least two annotations. We adopted a multi-label annotation setup. Annotators were presented one input sentence at a time, and were asked to select at least one out of our 9 frames. If no frame matched, the

⁷<https://annotation-app-sltly.onrender.com/>.

⁸This study was approved by Human Ethics Committee (Reference No. 2025-32561-67793-4). Annotators were compensated with vouchers over 50 USD, well above local minimum wage.

Label	α	Label	α
Economic	0.89	Political / Policy	0.78
Health / Safety	0.87	Cultural Identity	0.76
Legality / Crime	0.82	Other	0.66
Security / Defense	0.79	Fairness / Equality	0.65
Public Opinion	0.79	Morality	0.63

Table 2: Krippendorff’s α Per Label, Averaged Across Annotation Sessions and Sorted in Descending Order.

category “Other” should be selected.

Inter-Annotator Agreement Given our multi-label setup, we assessed inter-annotator agreement in two ways. Following prior work (Arora et al., 2025), we use label-wise Krippendorff’s α to assess per-label agreement, and the Jaccard index to measure the averaged overlap in label sets between annotators. The averaged Krippendorff’s α across all labels and annotators indicates moderate to high agreement ($\alpha = 0.76$), with some variance across frame labels (Table 2), however even the frame category with the lowest α reaches moderate agreement. We also compute an averaged Jaccard index of 0.79 which confirms this trend, suggesting a substantial overlap across label sets. This shows that a more robust set of labels applied at a sentence level allows to achieve substantial agreement despite the purported subjectivity of the task.

Gold Standard To construct a gold standard for model evaluation, we applied majority voting among annotators for each sentence. A clear consensus (where either two or three annotators chose the same label) was achieved for 532 cases (where in 500 cases a single label was assigned, while in 32 all annotators have chosen two consistent labels). Sentences where annotators disagreed (45) were adjudicated by two authors of the paper, who independently re-assessed the cases of contention. We kept only the cases where both adjudicators agreed with one of the original annotations, discarding 13 sentences. This process resulted in a validation data set of 556 sentences, which we use in Section 6.1 to assess model performance.

6 Methodology

Having established a general framework, and a conceptualization of media frames that generalize across articles and comments, we next describe the implementation of the framework components

RoBERTa			Llama			Qwen		
P	R	F1	P	R	F1	P	R	F1
0.70	0.65	0.66	0.68	0.65	0.65	0.64	0.64	0.64

Table 3: Precision (P), Recall (R) and Macro-F1 Scores for the Three Fine-Tuned LMs on MFC + Hartmann Data.

on our SOCC and NYT data sets of articles and comments (Section 4.1).

6.1 Frame Classification

We train frame classifiers that predict sentence-level media frames for both comments and articles, following the framing schema defined in Section 5.1. To do so, we require suitable training data from both document types, and a robust evaluation setup.

Training Data As our FrAC corpus is too small to support fine-tuning, we derive large-scale labeled data sets from the MFC (news articles) and Hartmann et al. (2019) (comments). We sentence-split both datasets and derive sentence-level labels from the original span-level labels: we first map the original frame labels to our 9 classes; then we identify spans with matching labels from at least two annotators with a minimum three-word overlap. Following previous methods (Naderi and Hirst, 2017; Hartmann et al., 2019; Park et al., 2022), agreed spans were mapped to the sentence-level: sentences fully covered by agreed spans inherited the span labels. This results in approximately 63,000 labeled sentences predominantly from the MFC (including 625 sentences from the much smaller Hartmann et al. (2019) dataset). Further details on frame and topic distributions of the training data are provided in Appendix D, Table 9.

Frame Classification Models To automatically classify frames in news articles and comments, we fine-tuned three models – RoBERTa-Large (Liu et al., 2019b) (full fine-tuning), LLaMA3.1-8B-Instruct (Grattafiori et al., 2024), and Qwen2.5-7B-Instruct (Qwen et al., 2025) (both parameter-efficient fine-tuning) – using the sentence-level annotated data described above with a random 80/10/10 (training/dev/test) split, stratified by frame and topic. Fine-tuning parameters are reported in Appendix A, and prompting details in Appendix B.

Label	P	R	F1
Economic	0.77	0.83	0.80
Morality	0.70	0.70	0.70
Fairness and Equality	0.68	0.59	0.64
Legality and Crime	0.83	0.86	0.85
Political and Policy	0.77	0.85	0.81
Security and Defense	0.69	0.59	0.63
Health and Safety	0.73	0.80	0.76
Cultural Identity	0.71	0.54	0.61
Public Opinion	0.61	0.61	0.61
Other	0.53	0.14	0.22

Table 4: Per-Frame Precision, Recall, and Macro-F1 Scores for Fine-Tuned RoBERTa on MFC + Hartmann Data.

X-Domain			FrAC (A)			FrAC (C)		
P	R	F1	P	R	F1	P	R	F1
0.65	0.59	0.59	0.74	0.81	0.77	0.84	0.85	0.83

Table 5: Precision, Recall and Macro-F1 Scores of RoBERTa Cross-Domain (X-Domain) and on FrAC Articles (A) and Comments (C).

Evaluation The averaged results in Table 3 show little performance difference across models, with RoBERTa-Large slightly outperforming the larger Qwen2.5-Instruct and LLaMA3.1-8B-Instruct. This result aligns with recent literature on frame classification, where BERT-based pre-trained models have consistently demonstrated strong performance (Sajwani et al., 2024; Daffara et al., 2025). Since RoBERTa also requires substantially less fine-tuning time (approximately 20 minutes versus 13 hours for PEFT-tuning of LLaMA and Qwen), all remaining experiments in this paper are based on fine-tuned RoBERTa. Table 4 reports the per-label breakdown for the best performing model, RoBERTa.

To evaluate model generalizability to unseen topics, we train RoBERTa on all but one MFC topics, evaluate on the held-out topic, averaged over all five combinations. Table 5 (X-Domain) revealed an expected, albeit relatively small drop in RoBERTa performance (from 0.66 In-Domain in Table 4 to 0.59), suggesting a reasonable ability to generalize to unseen topics, which is important because the model will be applied to topics not present in the MFC training set.

Finally, we evaluated RoBERTa on the annotated FrAC corpus (Table 5, FrAC), i.e., a data set that differs from the MFC fine-tuning data in its

outlets, time points and distinct yet overlapping set of topics (Section 5 and Table 6). The model achieved strong performance on FrAC with an F1 score of 0.77 on articles and 0.83 on comments, demonstrating effective generalization to reader comments despite being trained primarily on formal news articles.⁹ This result validates our central methodological contribution: that sentence-level frame classification can successfully bridge the gap between formal news articles and informal reader comments.

Having validated RoBERTa’s performance on our gold-standard FrAC corpus, we then applied it to the full sentence-tokenized SOCC and NYT data sets (Section 4.1) for our large-scale framing effects analysis.¹⁰

6.2 Primary Frame Detection

We analyze framing effects between articles and comments in terms of their respective primary frames as well as article secondary frames. We define the primary frame as the most frequent sentence-level frame label in a given document. Secondary frames are all other frames that were assigned to at least one sentence. To reduce noise from near ties, we only assign a primary frame if the most frequent frame meets both of the following criteria: it appears in at least 3 sentences and it represents at least 40% of all sentences in the document. We make an exception for single-sentence texts (comments), and consider their single frame automatically to be primary. Articles and comments whose most frequent frame fails to meet these combined thresholds were excluded from further analysis.

Evaluation To assess the reliability of our primary frame detection heuristics, two authors of the paper examined a sample of 110 articles and comments (55 for each) and assessed if the primary frame derived by applying the heuristics explained above is correct. The agreement rate between the annotators was 83%. Furthermore, human annotations agreed with the heuristic prediction 88% of the time, showing that our approach effectively

⁹FrAC was annotated directly at the sentence level using the consolidated 9+1 frame taxonomy, whereas MFC labels were created post-hoc from span-level annotations, introducing label noise which likely negatively affected model performance.

¹⁰We used the NLTK library (<https://www.nltk.org/>) for sentence tokenization of article and comment texts.

Topic	SOCC			NYT		
	Articles	Comments	Avg. C/A	Articles	Comments	Avg. C/A
Gun control	20	855	42	112	28,137	251
Russia-Ukraine	77	2,171	28	310	100,355	323
Trump & Elections	180	7,780	43	247	61,550	249
Healthcare	53	1,636	30	326	67,124	205
Immigration	115	4,394	38	172	33,060	192
LGBT+ Rights	23	803	34	22	3,279	149
Education	90	3,907	43	78	11,614	149
Abortion	17	508	29	38	6,193	163
Israel-Palestine	53	1,794	34	63	7,417	118
Climate Change	311	11,271	36	231	23,990	104
Syria & IS	138	4,195	30	72	3,757	52
Total	1,077	39,614	36.8	1,699	346,476	204

Table 6: SOCC and NYT Data Statistics by Topic and Source, including the average number of comments per article (Avg. C/A), after pre-processing as explained in Section 6.

captures the primary frame for both articles (95%) and comments (82%).

6.3 Topic Extraction

We apply BERTopic (Grootendorst, 2022) separately to SOCC and NYT to extract topics from news articles, assigning each article a probability distribution across all identified topics.¹¹ We analyzed the initial hierarchical topic structures produced by BERTopic and manually merged closely positioned topics with small inter-topic distances and high thematic similarity. The refined topic model was then applied to reader comments. We ultimately assigned each article and each comment the single most likely topic predicted by BERTopic.

This process was conducted independently for both corpora, to avoid any bias in subsequent analyses. The extracted topics showed substantial thematic overlap across data sets, enabling cross-dataset comparison. We manually identified eleven overarching, semantically cohesive topics which we use in our downstream analysis (Table 6).

Evaluation To assess the quality of topic assignments, we randomly sampled 5 articles and 5 comments from each of the 11 topic groups, yielding a total of 110 texts. Two authors of the paper independently reviewed each assigned topic and judged whether it was acceptable or not. We observe a high degree of inter-annotator agreement (96%). The overall accuracy of the model’s predicted topics, calculated against the labels assigned by the

annotators was 93% (on both comments and articles).

6.4 Resulting Data Set

With our framework and validated implementation at hand, we next apply it to analyze framing effects across the SOCC and NYT data sets (Section 4.1) and eleven selected topics (Section 6.3). We only include article-comment pairs where both share the same predicted topic (excluding off-topic comments to study framing effects when discussions remain on-topic) and where primary frames could be extracted for both articles and comments. We also filter out comments shorter than 5 words, and duplicate comments. This final data set consists of nearly 3,000 articles and over 380,000 associated comments, as detailed in Table 6.

7 Framing Effects across Topics and Outlets

We now address our three research questions (Section 3) asking to what extent primary article frames are taken up in readers’ comments? How do these trends vary by topic? And, if the primary frame does not stick, how readers reconstruct the media frame in their comments?

7.1 To what extent are primary article frames retained in reader comments?

In line with our definition of frame retainment as the echoing of the *primary* frame of a news article, we first investigate how often the primary frame kept in its associated direct readers comments. We

¹¹We report the full list of parameters used in Appendix C.




Figure 2: Predicted Frame Retention Rates in Reader Comments based on Regression-Derived Marginal Effects. X- and Y-axes show predicted retention for SOCC and NYT, respectively. Retention varies by frame type, and the predictions correlate strongly across the two media outlets.




Figure 3: Predicted Topic Retention Rates in Reader Comments based on Regression-Derived Marginal Effects. Points above the diagonal line indicate higher predicted retention in NYT, while points below indicate higher predicted retention in SOCC. Topic names are shortened for readability.

do so across the New York Times (NYT; USA) and The Globe and Mail (SOCC; Canada).

Overall, we find that NYT readers demonstrate stronger frame retention (50.7%) compared to SOCC readers (37.3%). Conversely, at least half of the comments (almost two thirds in SOCC) *change* away from the primary article frame indicating that readers often do not retain the original primary framing of the article. In section 7.3, we analyze these reframing patterns in depth.

We next investigate how this pattern varies by frame type. We do so through a logistic regression mixed effects model that predicts frame retention (i.e., comment frame being identical to the primary article frame) as a function of fixed effects for primary article frame and article topic, and a random effect controlling for article ID,¹²

$$y_i \sim \text{topic}_d + \text{frame}_d + (1|\text{id}_d), \quad (1)$$

where the input is a document-response pair $i = (d, r)$, y_i is a binary indicator of frame retention ($\text{frame}_d = \text{frame}_r$).

Here, we are interested in the marginal effect of frame_d on y as derived from the fitted model coefficients. Figure 2 displays the frame retention patterns separately for each primary article frame.

¹²We control for article ID due to repeated measurements, given that we have multiple comments per article.

Looking at the spread of frames along the x- and y-axis independently, we can see that the extent of retention varies by frame. Specifically, frames such as “Political and Policies”, “Economic”, and “Health and Safety” show high retention rates. In contrast, “Public Opinion”, “Fairness and Equality”, and “Cultural Identity” exhibit substantially lower retention rates. This suggests that readers are more likely to reframe content that is presented through frames that require alignment with personal beliefs or values to be accepted, in line with theoretical postulations in framing research (Chong and Druckman, 2007).

To formally test whether the frame choices in reader comments were influenced by the frames used in the corresponding articles – that is, whether the distribution of frames in comments is not independent of the article’s primary frame – we conducted Chi-squared tests of independence. For each source frame f , we tested the null hypothesis that the presence of f in an article is independent of the presence of f in a comment. For all frames, we were able to reject the null hypothesis of independence ($p \ll .001$), although with small to moderate effect sizes (e.g., small effects: “Public Opinion” Cramér’s $V = 0.065$ in NYT; moderate effects: “Economic” $V = 0.357$ in SOCC).

Figure 2 additionally reveals that the reframing patterns are remarkably consistent across NYT and

SOCC, i.e., all points cluster close to the diagonal line of perfect correlation. This cross-outlet consistency addresses a key limitation of previous single-outlet studies and suggests that our uncovered patterns might reveal a general trend in news reader behavior, rather than being idiosyncratic to a particular outlet or sample.

7.2 How does frame retention vary across topics?

While previous studies have exclusively focused on a single topic and have often produced contradictory results, here we can more rigorously inspect the variation of framing effects across topics under our unified framework.

We explore the effect of topics on framing effects using the same regression framework as in Section 7.1, but this time focusing on the marginal effects of topic on retention from equation (1). This provides predicted frame retention probabilities per topic, controlling for frame type and article-specific effects. The resulting predicted probabilities of frame retention per topic are shown in Figure 3 for NYT (y-axis) and SOCC (x-axis).

We find that predicted framing effects vary considerably by topic, with notable differences between outlets. Several topics show consistent retention patterns across both outlets. Trump & Elections, for instance, demonstrates high predicted retention in both outlets. In contrast, international conflicts like Israel-Palestine and Syria & IS show consistently low retention across both outlets, clustering in the lower portion of the distribution.

Other topics reveal substantial outlet-specific differences. Immigration and Healthcare show notably higher predicted retention in NYT than SOCC, as does coverage of the Russia-Ukraine conflict and LGBT+ Rights, as shown in Figure 3.

These outlet-specific variations, which have not been examined in previous works, confirm that “different kinds of issues are interpreted by the media and by the public in different ways” (Neuman et al., 1992, pg. 17), suggesting that topic-based retention and reframing patterns are more outlet-specific than the frame-based patterns in Figure 2. This finding is in line with the inconsistent results reported in previous small-scale, single-topic studies in the social sciences (Sections 2 and 3).

In this paper, we deliberately remain on the *descriptive* level and refrain from *explaining* the observed differences, being cognizant of many con-

foundings factors that may be at play. However, our framework paves the way for future causal analyses of framing across different topics and news outlets.

7.3 If readers reject the primary article frame, how do they reframe?

Prior communication studies have used varied approaches to analyze reframing - from binary annotations of the presence of frames (e.g., Holton et al. (2014); Muñiz et al. (2015)) to primary frame identification (e.g., Zhou and Moy (2007); Koteyko et al. (2013)). However, none systematically examined *how* readers reframe when they diverge from the primary article frame - whether they introduce entirely novel frames or selectively emphasize non-primary frames already present in the source. Here, we address this question directly.

Comments change away from the primary article frame 50% (NYT) or 63% (SOCC) of the time, so a natural question is *how* readers reframe articles. In our analysis, we distinguish between two types of reframing: (a) **complete reframing**, where the comment frame does not appear in the source article at all, and (b) **selective reframing**, where comments adopt secondary frames which are present in the source article, but do not dominate the message. In SOCC, 11.7% of comment frames exhibit complete reframing, while 62.8% show selective reframing. In NYT, these fractions are 5.2% and 49.3% respectively. This shows that, rather than completely accepting or rejecting the primary source framing, audiences predominantly engage in selective interpretation by adopting secondary frames from the article.

Figure 4 shows reframing patterns for the NYT. Focusing first on complete reframing (Figure 4a), Political framing dominates source articles but commentators introduce diverse alternative frames. In particular, readers reframe in terms of value-laden frames including “Morality” (increasing from 0.8% in articles to 26.4% in Comments) or “Cultural Identity”, as well as “Economic” framing. This pattern holds across both NYT and SOCC (as shown in Appendix G).

Selective reframing patterns in NYT are shown in Figure 4b. Here, the frame distributions in the comments align more closely with the primary article frame distributions, although we again observe an increase in value-laden frames (“Morality”, “Health and Safety”) as well as “Economic” framing.




Figure 4: Reframing in the New York Times. (a) Complete reframing patterns where readers use frames that never appear anywhere in the original articles. (b) Selective reframing where readers emphasize secondary frames that differ from the primary article frame but occur in the article. Left percentages show primary article frame distributions; right percentages show comment frame distributions. Left-side blocks (article frames) are not scaled to proportion in order to avoid visual clutter.

Topic-specific Reframing We investigate topic-specific reframing patterns for topics that included at least 100 articles. For instance, in the case of *Immigration* (Table 7, left), SOCC readers engage in selective reframing by emphasizing “Cultural Identity” frames when “Political and Policies” dominate the article, and vice versa, suggesting a tight interconnection of these two perspectives. NYT readers show different selective reframing patterns, primarily shifting from “Legality and Crime” frames to “Political and Policies” frames, and from “Political and Policies” to “Economic” frames, indicating a tendency toward more policy-oriented and economic interpretations of immigration coverage. Complete reframing differs substantially (Table 7, right). In NYT, complete reframing shows “Economic” frames being introduced when they were entirely absent from articles. Again, both outlets reframe more objective (Economic, Political, Legal) views into value-laden frames (Cultural, Morality, Fairness). Tables 11 and 12 in the Appendix show selective and complete reframing analysis for additional topics, revealing substantial topic-specific differences.

In conclusion, the results discussed here align with the constructionist view of framing as “interactive, vulnerable and prone to counter-frames” (Gorp, 2007, pg. 76), with audiences engaging in topic-specific patterns of selective retention for frames that resonate with them, whilst actively reframing those that do not.

	Selective Reframing	Complete Reframing
SOCC	Cultural → Political	Political → Cultural
	Health → Cultural	Economic → Political
	Political → Cultural	Health → Morality
	Health → Political	Economic → Cultural
	Economic → Political	Political → Fairness
NYT	Legality → Political	Political → Economic
	Legality → Economic	Political → Legality
	Political → Economic	Legality → Cultural
	Political → Legality	Political → Cultural
	Political → Cultural	Political → Morality

Table 7: Reframing “Immigration” in SOCC (Top) and NYT (Bottom). Each cell shows the five most common selective (left) and complete (right) shifts from primary article → comment frame. Frame names are abbreviated to their first word.

8 Discussion

Unlike previous computational approaches to framing detection and analysis, our study is centered on a *constructionist* understanding of framing as an interactive process where readers selectively engage, modify, and reconstruct media frames based on their own interpretative resources (Neuman et al., 1992; Scheufele, 1999; Gorp, 2007). While existing computational methods treat framing as isolated classification tasks focused solely on source content, we examined the dynamic interaction between media framing and audience reception. In uncovering systematicities (RQ1; Figure 2) and topic-specific differences (RQs 2-3) in a large-scale

study, we bridged the disconnect between large-scale yet isolated NLP approaches and small-scale communication studies that lack generalizability.

In applying one unified framing conceptualization across eleven topics, our findings suggest that inconsistent results on reframing patterns in prior work can be partially explained by topic- and outlet-specific variation in framing effects. We also identified consistent trends, such as value-laden frames (“Fairness and Equality”, “Morality”) being more likely to be reframed (Figure 2), but also more likely to be introduced in the context of complete reframing (Figure 4a). This divergent pattern suggests that audiences reject value-laden framings that conflict with their worldview while actively introducing those that align with it. This extends theoretical formulations about audience personal formulation of messages (Chong and Druckman, 2007) by showing that readers not only selectively accept frames matching their values, but also actively impose their preferred value frames onto discussions where they were originally absent.

In distinguishing selective from complete reframing, we show that readers predominantly reframe in selective ways, providing empirical evidence for what framing theory has long posited: that audiences reformulate and engage with content selectively based on their preferences (Neuman et al., 1992; Scheufele, 1999) – a phenomenon that small-scale social science studies could not disentangle systematically, and that NLP studies have overlooked by disregarding response analysis.

Our analysis reveals that frame-based patterns are relatively stable across outlets, while topic-specific effects vary more widely. These differences likely arise from complex interactions between unobservable factors including reader demographics and national context. Our data sets conflate country (US versus Canada), editorial leaning (center-left NYT versus center-right SOCC), and time period (2012–16 NYT vs 2017–18 SOCC). We hence maintain a *descriptive* level and leave disentanglement of these factors for future work.

A key methodological contribution lies in reliable frame classifiers across different text genres – formal news articles and informal reader comments. This cross-genre generalization addresses a critical gap in computational framing analysis, where models typically operate within homogeneous text types. Our validation shows strong performance on our human-annotated data set FrAC.

Our study builds on the MFC, a widely-used framework in NLP. Yet, it is by no means the only or optimal framework, and it has been criticized for the broad nature and lack of theoretical grounding (Ali and Hassan, 2022). We note that our method is framework-agnostic: alternative framing theories could be readily integrated, whether more generic frameworks (e.g., Semetko and Valkenburg (2000)), or issue-specific taxonomies tailored to particular domains. Similarly, our approach supports both large-scale analyses of framing trends and more fine-grained investigations. Future work could further explore the role of secondary frames.

Finally, our framework holds potential relevance for other NLP tasks including argumentation interpretation and generation, and disinformation detection. The systematic patterns we identify in how audiences reconstruct frames could inform models of persuasive text generation, helping systems understand how different framings might be received and transformed by target audiences. Other open questions for future work include understanding how framing effects propagate through nested comment conversations, identifying causal linguistic properties that explain specific framing effects, and conducting thread and temporal-level analyses.

9 Conclusion

This work addresses the gap between framing theory and its computational operationalization. While framing theory posits that individuals respond differentially to framed information – accepting, rejecting, or modifying frames based on their cognitive schemas – computational approaches have largely reduced framing to taxonomic classification, neglecting audience reception processes.

We introduced a scalable computational framework for analyzing media framing effects as an interactive process between source content and audience responses. Through analysis of over 3,000 news articles and 380,000 comments across two outlets and eleven topics, we demonstrate that readers reframe content at least as frequently as they retain original frames, often reframing selectively by picking up on secondary article frames. Frame retention varies systematically by frame type with objective frames (e.g., “Health & Safety”) more likely to be retained than value-laden frames (e.g., “Morality”). We contributed a robust framing scheme with classifiers that work across document types, providing a foundation for future research.

Acknowledgment

This paper was written with the support from the Melbourne Research Scholarship by the University of Melbourne provided to MG. This work was supported by the Australian Research Council Discovery Early Career Research Award (Grant No. DE230100761). We thank Damian Curran, Thomas Hollow, Jemima Kang, Tom Monaghan, Gordon William Smith, and Katie Warburton for their annotations.

References

- Afra Feyza Akyürek, Lei Guo, Randa Elanwar, Prakash Ishwar, Margrit Betke, and Derry Tanti Wijaya. 2020. [Multi-label and multilingual news framing analysis](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8614–8624. Association for Computational Linguistics.
- Mohammad Ali and Naeemul Hassan. 2022. [A survey of computational framing analysis approaches](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9335–9348, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Arnav Arora, Srishti Yadav, Maria Antoniak, Serge Belongie, and Isabelle Augenstein. 2025. [Multi-modal framing analysis of news](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 31519–31541, Suzhou, China. Association for Computational Linguistics.
- Amber E. Boydston, Dallas Card, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2014. [Tracking the development of media frames within and across policy issues](#). In *APSA 2014 Annual Meeting Paper*.
- 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 *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 438–444, Beijing, China. Association for Computational Linguistics.
- Dennis Chong and James Druckman. 2007. [Framing theory](#). *Annual Review of Political Science*, 10.
- Agnese Daffara, Sourabh Dattawad, Sebastian Padó, and Tanise Ceron. 2025. [Generalizability of media frames: Corpus creation and analysis across countries](#). In *Proceedings of the 14th Joint Conference on Lexical and Computational Semantics (*SEM 2025)*, pages 83–99, Suzhou, China. Association for Computational Linguistics.
- Lauren Michelle Dargay. 2016. [Relationships between elite news frames and frames in user comments: An analysis of terrorism coverage and follow-up comments on the New York Times online](#). MA Thesis, Kent State University, College of Communication and Information, Kent, Ohio. Advisors: Chance York, Gary Hanson; Committee Member: Michael Beam.
- Paul DiMaggio, Manish Nag, and David Blei. 2013. [Exploiting affinities between topic modeling and the sociological perspective on culture: Application to newspaper coverage of U.S. government arts funding](#). *Poetics*, 41(6):570–606.
- Robert Entman. 1993. [Framing: Toward clarification of a fractured paradigm](#). *The Journal of Communication*, 43:51–58.
- Anjalie Field, Doron Kliger, Shuly Wintner, Jennifer Pan, Dan Jurafsky, and Yulia Tsvetkov. 2018. [Framing and agenda-setting in Russian news: A computational analysis of intricate political strategies](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3570–3580, Brussels, Belgium. Association for Computational Linguistics.
- Baldwin Van Gorp. 2007. [The constructionist approach to framing: Bringing culture back in](#). *Journal of Communication*, 57(1):60–78.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

- Maarten Grootendorst. 2022. [Bertopic: Neural topic modeling with a class-based tf-idf procedure](#). *arXiv preprint arXiv:2203.05794*.
- Mareike Hartmann, Tallulah Jansen, Isabelle Augenstein, and Anders Søgaard. 2019. [Issue framing in online discussion fora](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1401–1407, Minneapolis, Minnesota. Association for Computational Linguistics.
- Avery Holton, Nayeon Lee, and Renita Coleman. 2014. [Commenting on health: A framing analysis of user comments in response to health articles online](#). *Journal of Health Communication*, 19(7):825–837. Publisher: Taylor & Francis.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. LoRA: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Austin Hubner and Graham Dixon. 2023. Is news media sharing an active framing process? Examining whether individual tweets retain news media frames about climate change. *Human Communication Research*, 49(1):75–84.
- Yangfeng Ji and Noah A. Smith. 2017. [Neural discourse structure for text categorization](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 996–1005, Vancouver, Canada. Association for Computational Linguistics.
- Kristen Johnson and Dan Goldwasser. 2018. [Classification of moral foundations in microblog political discourse](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 720–730, Melbourne, Australia. Association for Computational Linguistics.
- Shima Khanehzar, Trevor Cohn, Gosia Mikolajczak, Andrew Turpin, and Lea Frermann. 2021. [Framing unpacked: A semi-supervised interpretable multi-view model of media frames](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for*
- Computational Linguistics: Human Language Technologies*, pages 2154–2166. Association for Computational Linguistics.
- Shima Khanehzar, Andrew Turpin, and Gosia Mikolajczak. 2019. [Modeling political framing across policy issues and contexts](#). In *Proceedings of the 17th Annual Workshop of the Australasian Language Technology Association*, pages 61–66, Sydney, Australia. Australasian Language Technology Association.
- Jelena Kleut and Ana Milojevic. 2021. [Framing protest in online news and readers’ comments: The case of Serbian protest "Against Dictatorship"](#). *International Journal of Communication*, 15(0):21. Number: 0.
- Veselin Kolhatkar, Haofen Wu, Liane Cavasso, and Maite Taboada. 2020. [The SFU opinion and comments corpus: A corpus for the analysis of online news comments](#). *Corpus Pragmatics*, 4(2):155–190. Published: 2 November 2019; Issue Date: June 2020.
- Nelya Koteyko, Rusi Jaspal, and Brigitte Nerlich. 2013. [Climate change and ‘climategate’ in online reader comments: A mixed methods study](#). *The Geographical Journal*, 179(1):74–86.
- Siyi Liu, Lei Guo, Kate Mays, Margrit Betke, and Derry Tanti Wijaya. 2019a. [Detecting frames in news headlines and its application to analyzing news framing trends surrounding U.S. gun violence](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 504–514, Hong Kong, China. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Julia Mendelsohn, Ceren Budak, and David Jurgens. 2021. [Modeling framing in immigration discourse on social media](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2219–2263. Association for Computational Linguistics.

- Carlos Muñiz, Salvador Alvidrez, and Nilsa Téllez. 2015. Shaping the online public debate: The relationship between the news framing of the expropriation of YPF and readers' comments. *International Journal of Communication*, 9:3245–3263.
- Nona Naderi and Graeme Hirst. 2017. [Classifying frames at the sentence level in news articles](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 536–542, Varna, Bulgaria. INCOMA Ltd.
- Thomas E. Nelson and Zoe M. Oxley. 1999. [Issue framing effects on belief importance and opinion](#). *The Journal of Politics*, 61(4):1040–1067.
- W. Russell Neuman, Marion R. Just, and Ann N. Crigler. 1992. *Common Knowledge: News and the Construction of Political Meaning*. University of Chicago Press, Chicago.
- Yulia Otmakhova, Shima Khanehzar, and Lea Frermann. 2024. [Media framing: A typology and survey of computational approaches across disciplines](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15407–15428, Bangkok, Thailand. Association for Computational Linguistics.
- Chan Young Park, Julia Mendelsohn, Anjalie Field, and Yulia Tsvetkov. 2022. [Challenges and opportunities in information manipulation detection: An examination of wartime Russian media](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5209–5235, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jakub Piskorski, Nicolas Stefanovitch, Nikolaos Nikolaidis, Giovanni Da San Martino, and Preslav Nakov. 2023. [Multilingual multifaceted understanding of online news in terms of genre, framing, and persuasion techniques](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3001–3022, Toronto, Canada. Association for Computational Linguistics.
- Vincent Price, David Tewksbury, and Elizabeth Powers. 1997. Switching trains of thought: The impact of news frames on readers' cognitive responses. *Communication Research*, 24(5):481–506.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025. [Qwen2.5 technical report](#). *arXiv preprint arXiv:2412.15115*.
- Julian Risch and Ralf Krestel. 2020. [A dataset of journalists' interactions with their readership: When should article authors reply to reader comments?](#) In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management, CIKM '20*, page 3117–3124, New York, NY, USA. Association for Computing Machinery.
- Shamik Roy and Dan Goldwasser. 2020. [Weakly supervised learning of nuanced frames for analyzing polarization in news media](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7698–7716. Association for Computational Linguistics.
- Ahmed Sajwani, Alaa El Setohy, Ali Mekky, Diana Turmakhan, Lara Hassan, Mohamed El Zeftawy, Omar El Herraoui, Osama Mohammed Afzal, Qisheng Liao, Tarek Mahmoud, Zain Muhammad Mujahid, Muhammad Umar Salman, Muhammad Arslan Manzoor, Massa Baali, Jakub Piskorski, Nicolas Stefanovitch, Giovanni Da San Martino, and Preslav Nakov. 2024. [FRAPPE: FRAMing, persuasion, and propaganda explorer](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 207–213, St. Julians, Malta. Association for Computational Linguistics.
- Dietram A. Scheufele. 1999. [Framing as a theory of media effects](#). *Journal of Communication*, 49(1):103–122.

- Holli A Semetko and Patti M Valkenburg. 2000. Framing European politics: A content analysis of press and television news. *Journal of Communication*, 50(2):93–109.
- Fuyuan Shen. 2004. Effects of news frames and schemas on individuals' issue interpretations and attitudes. *Journalism & Mass Communication Quarterly*, 81(2):400–416.
- Patti Valkenburg, Holli Semetko, and Hubert Claes. 1999. [The effects of news frames on readers' thoughts and recall](#). *Communication Research*, 26:550–569.
- Yuqiong Zhou and Patricia Moy. 2007. [Parsing framing processes: The interplay between online public opinion and media coverage](#). *Journal of Communication*, 57(1):79–98.

A Frame Classifiers

We report the parameters used for fine-tuning the two LLMs — Llama and Qwen — and RoBERTa for the frame classification task. The LLMs were fine-tuned using Unsloth¹³, applying parameter-efficient fine-tuning via Low-Rank Adaptation (LoRA) (Hu et al., 2022), while RoBERTa was fully fine-tuned. All training was conducted on high-performance computing infrastructure with NVIDIA A100 GPUs. We also include the final prompt used for the task in Table 8.

A.1 LLaMA3-8B-Instruct and Qwen2.5-Instruct

We used the same set of parameters on Unsloth for both models.

- Max Sequence Length: 3000 tokens
- Batch Size: 8 (gradient accumulation = 8, effective batch = 64)
- Number of Epochs: 6
- Learning Rate: 2e-4
- LoRA Configuration: $r = 16$, $\alpha = 16$, dropout = 0, target modules = [q_proj, k_proj, v_proj, o_proj, gate_proj, up_proj, down_proj]
- Precision: bf16 (if supported) or fp16
- Optimizer: AdamW (8-bit)
- Scheduler: Linear with 10% warm-up ratio
- Regularization: Weight decay = 0.01
- Early Stopping: Patience = 3 epochs
- Checkpointing: Save best model by validation loss at end of each epoch
- Packing: Enabled

A.2 RoBERTa

- Maximum sequence length: 70 tokens
- Data splitting:
 - nosep Training set: 80%
 - nosep Validation set: 10%
 - nosep Test set: 10%
- Batch size: 64
- Learning rate: 5e-6
- Number of epochs: 10
- Weight decay: 0.01
- Optimizer scheduler: Linear learning rate scheduler with 5% warmup
- Early stopping: Patience of 2 epochs based on validation F1 score
- Evaluation and model saving: After every epoch, best model saved based on F1 score
- Loss function: Cross-entropy

B LLMs Prompt

We experimented with various prompting strategies to balance performance and prompt length. We found that including an elaborated description of the category definitions from Card et al. (2015), along with clear task instructions, consistently yielded the best results. The final prompt is reported in Table 8.

C BERTopic Configuration

We trained a BERTopic model (Grootendorst, 2022) for clustering textual topics with the following configuration:

- **Embedding Model:** all-mpnet-base-v2 from the SentenceTransformers library.
- **Vectorizer:** CountVectorizer with the following parameters:
 - stop_words: English stopwords from NLTK
 - min_df: 5

¹³<https://unsloth.ai/>

Prompt Used for Fine-Tuning You are an expert in discourse and framing analysis. You are analyzing how texts are framed in their coverage of topics. Your task is to identify the frames used in the text into one of nine frame categories based on how the issue is presented and which aspects are emphasized. Instructions: 1. Read the text carefully. 2. Identify the frame that shapes the article’s narrative. Only ONE frame applies. 3. Output ONLY the NUMERICAL code (e.g. "2") for those frames first, followed by a new line. Do not include any other text or explanation. 1 - Economic: The costs, benefits, or financial implications of the issue (individual, community, or economy) 2 - Morality: Perspectives compelled by religion, ethics, or sense of social responsibility 3 - Fairness and Equality: Equality or inequality in how laws and resources are applied or distributed 4 - Legality and Crime: Laws, judicial interpretations, crime rates, legal consequences 5 - Health and Safety: Healthcare, disease, sanitation, mental health, infrastructure safety 6 - Cultural Identity: Social norms, values, customs, and cultural preservation 7 - Public Opinion: Public attitudes, polling, consequences of diverging from popular opinion 8 - Security and Defense: Protection, national security, border security, social stability 9 - Political and Policies: Political discourse, partisanship, lobbying, legislative processes 10 - Other: If the sentence is not framed in any of the proposed categories, select this label

Table 8: Prompt Template Format Used for Fine-Tuning the Models.

- ngram_range: (1, 2)
- **UMAP Dimensionality Reduction:**
 - n_neighbors: 15
 - n_components: 5
 - min_dist: 0.0
 - metric: cosine
 - random_state: 42
- **Clustering Model:** HDBSCAN with:
 - min_cluster_size: 15
 - min_samples: 7
 - metric: euclidean
 - cluster_selection_method: eom
 - prediction_data: True
- **BERTopic Parameters:**
 - calculate_probabilities: True
 - verbose: True
 - n_gram_range: (1, 2)
 - min_topic_size: 5

D Processed Training Data

We report the statistics of the final processed data set, as explained in Section 6.1. Table 9 reports the number of final samples per label across topics, from both the MFC and Hartmann et al. (2019).

E Annotation Instructions and Interface

This section includes the instructions which were presented to each annotator on the web interface developed for the task (a screenshot of the interface is reported in Figure 5).

Understanding the Influence of News Framing on Readers

Hi there! Thanks for agreeing to participate in this annotation task.

If you’re viewing this on a phone or tablet, we recommend switching to a laptop for better comfort and

Frame	Guns	Death Penalty	Immig.	Same-Sex	Tobacco	Hartmann	Total
Legality and Crime	4,942	9,215	4,111	3,306	1,376	344	23,294
Political and Policies	7,405	1,006	3,436	3,357	1,853	211	17,268
Security and Defense	947	163	770	11	-	-	1,891
Health and Safety	938	740	534	158	1,409	-	3,779
Economic	810	156	1,294	441	2,303	70	5,074
Morality	183	573	147	999	64	-	1,966
Cultural Identity	636	432	722	552	459	-	2,801
Public Opinion	1,089	161	470	967	130	-	2,817
Fairness and Equality	60	677	284	681	34	-	1,736
Other	275	451	253	1,716	305	-	3,000
Total	17,285	13,574	12,021	12,188	7,933	625	63,626

Table 9: Distribution of Frame Categories Across Issue Topics in the Training Data Used for Fine-Tuning the Frame Classification Models. The data combines the Media Frames Corpus with additional annotations from [Hartmann et al. \(2019\)](#).

Figure 5: An Example From the Annotation Task on the Deployed Interface.

readability.

In this project, we’re investigating the different formulations of news in terms of frames.

What Are Frames?

Frames are the perspectives or angles through which information is presented in media. They highlight certain aspects of an issue while downplaying others, influencing how audiences understand and respond to news.

Think of frames as different “lenses” that focus attention on specific elements of a story while leaving others in the background.

Frames answer the question: *“How is the issue at hand being presented?”*

Your Task

Your task is to identify which frames are used in short snippets of text from various sources. This helps us understand how different issues are presented and framed in different contexts.

How the Annotation Works

You will see one text at a time. Please read each text carefully and choose the frame that best represents how the information is being conveyed by clicking on the corresponding label (colourful boxes).

If you think more than one frame applies, please select all that you believe are relevant.

If none of the provided categories apply, please select the label “Other.”

Please note: the examples all focus on same-sex marriage for clarity. The real annotation task covers a variety of different issues.

Once you start annotating, you can click on the “Instructions” button at any time to retrieve this information.

Framing Categories

See Table 10

Frame	Description	Example
Economic	The costs, benefits, or financial implications of the issue (individual, community, or economy).	<i>Legalizing same-sex marriage would boost state economies through increased spending on weddings and tourism.</i>
Morality	Perspectives compelled by religion, ethics, or sense of social responsibility.	<i>Same-sex marriage is morally wrong and goes against the teachings of our faith.</i>
Fairness and Equality	Equality or inequality in how laws and resources are applied or distributed.	<i>Denying same-sex couples the right to marry is a clear violation of equal protection under the law.</i>
Legality and Crime	Laws, judicial interpretations, crime rates, legal consequences.	<i>The Supreme Court ruled that bans on same-sex marriage are unconstitutional and violate civil liberties.</i>
Health and Safety	Healthcare, disease, sanitation, mental health, infrastructure safety.	<i>Allowing same-sex marriage is an endorsement of unsafe sexual practices and will contribute to the spread of disease.</i>
Cultural Identity	Social norms, values, customs, and cultural preservation.	<i>Same-sex marriage threatens the traditional concept of family that our culture has upheld for generations.</i>
Public Opinion	Public attitudes, polling, consequences of diverging from popular opinion.	<i>A recent poll shows that 65% of Americans now support same-sex marriage, up from 30% a decade ago.</i>
Security and Defense	Protection, national security, border security, social stability.	<i>Same-sex marriage could undermine traditional social structures and thus contribute to community stability and security.</i>
Political and Policies	Political discourse, partisanship, lobbying, legislative processes.	<i>The governor vetoed a bill legalizing same-sex marriage, citing pressure from conservative lawmakers.</i>
Other	If the sentence is not framed in any of the proposed categories, select this label.	—

Table 10: Framing Categories Shown on the Annotation Interface.

Practice Phase

Before starting the actual annotation task, we’ll practice together with three examples.

During the practice phase:

- You will be shown **one sentence at a time**.
- You should **select the appropriate frame label(s)** by clicking on the colored boxes.
- After you make your selection, the **correct answer will appear underneath with an explanation**.

Take your time to read each sentence carefully and consider which perspective or angle is being emphasized.

You can consult the examples you will see during the practice at any time during the annotation by clicking on the “Tutorial Example” button on the top-right.

Final Notes

If you have any questions, please contact the organizer.

For more information regarding the framing labels and theory, please refer to the [Policy Frames Codebook](#) and this [paper](#).

Happy annotating!

F Reframing Patterns by Topic

Figure 11 and Figure 12 show the 5 most frequent selective and complete reframing patterns in both news outlets in coverage of climate change, immigration, Syria & IS, Trump & elections.

Topic	Outlet	Selective Reframing Pattern	Count
CC	SOCC	Political → Economic	1001
		Economic → Political	854
		Political → Public Opinion	113
		Economic → Cultural	90
		Political → Cultural	79
	NYT	Political → Economic	1585
		Economic → Political	857
		Political → Health	428
		Political → Cultural	345
		Political → Legality	232
IM	SOCC	Cultural → Political	132
		Health → Cultural	113
		Political → Cultural	99
		Health → Political	96
		Economic → Political	85
	NYT	Legality → Political	1388
		Legality → Economic	955
		Political → Economic	927
		Political → Legality	903
		Political → Cultural	769
SY	SOCC	Political → Security	352
		Security → Political	326
		Security → Morality	78
		Security → Cultural	64
		Political → Morality	41
	NYT	Security → Political	321
		Health → Political	146
		Health → Morality	123
		Security → Morality	117
		Security → Cultural	79
TR	SOCC	Political → Cultural	319
		Political → Public Opinion	256
		Political → Economic	146
		Cultural → Political	145
		Political → Fairness	108
	NYT	Political → Cultural	2136
		Political → Morality	2070
		Cultural → Political	1145
		Political → Health	983
		Political → Economic	774

Table 11: Top 5 Selective Reframing Patterns by Topic and Outlet. CC = Climate Change, IM = Immigration, SY = Syria, TR = Trump.

G Reframing in SOCC

Topic	Outlet	Complete Reframing Pattern	Count
CC	SOCC	Political → Economic	125
		Political → Cultural	106
		Political → Health	72
		Political → Public Opinion	68
		Economic → Cultural	66
	NYT	Political → Morality	473
		Political → Health	258
		Political → Cultural	120
		Health → Morality	62
		Economic → Morality	51
IM	SOCC	Political → Cultural	32
		Economic → Political	23
		Health → Morality	22
		Economic → Cultural	17
		Political → Fairness	17
	NYT	Political → Economic	384
		Political → Legality	236
		Legality → Cultural	189
		Political → Cultural	181
		Political → Morality	177
SY	SOCC	Security → Political	79
		Security → Cultural	56
		Political → Morality	55
		Security → Morality	49
		Political → Cultural	39
	NYT	Security → Morality	48
		Security → Cultural	27
		Security → Economic	25
		Fairness → Morality	14
		Health → Cultural	14
TR	SOCC	Political → Cultural	88
		Political → Fairness	70
		Political → Economic	60
		Political → Legality	60
		Political → Morality	38
	NYT	Political → Morality	446
		Political → Health	379
		Political → Economic	265
		Political → Legality	252
		Political → Cultural	237

Table 12: Top 5 Complete Reframing Patterns by Topic and Outlet. CC = Climate Change, IM = Immigration, SY = Syria, TR = Trump.





Figure 6: Reframing patterns in SOCC. **Top: Complete** reframing showing cases where readers use frames never appearing in articles. **Bottom: Selective** reframing showing cases where readers emphasize secondary frames differing from the primary article frame but appearing elsewhere in articles. Left-side blocks (article frames) are not scaled to proportion to avoid visual clutter.