“my stance decides my language”: Modeling of Framing and Political Stance in News Media

Anonymous ACL submission

Abstract

Framing is a political strategy in which journalists and politicians highlight certain aspects of an issue or a problem to influence public opinion. Frameworks for detecting framing in news articles or social media posts are necessary in order to understand the spread of biased information in our society. Prior research efforts have shown that their framework for framing detection works well by predicting political affiliation afterward. In this paper, rather than predicting stance after detecting frames, we incorporate stance prediction into a framing detection model to jointly capture framing languages better. We take advantage of political stance data, which are more readily available than framing data that require manual annotation of professionals, and propose automatic framing detection models, which can detect previously unseen framing phrases. We compare two different methods of incorporation and show that leveraging stance prediction improves the separation of liberal and conservative biased frame language.

1 Introduction

Framing in social sciences refers to emphasizing desired aspects of an issue to promote a particular perspective (Entman, 1993). By selecting certain information and hence elevating the salience of that information, topics can be expressed with different frames. Research on frames has largely focused on political and social issues, such as the stances of politicians (Johnson and Goldwasser, 2016), the U.S. anti-nuclear war movement (Entman and Rojecki, 1993), stem cell research (Nisbet et al., 2003), and COVID-19 (Wicke and Bolognesi, 2020).

Detecting and analyzing framing is crucial in comprehending public perspectives and biases in social issues. In a world where people are overwhelmed with information from news media outlets and social media platforms, the importance of understanding framing cannot be overstated.

In response to the success of machine learning (ML), ML techniques have been applied to detect frames (Card et al., 2015; Guo et al., 2016; Johnson et al., 2017a; Bhatia et al., 2021). In many framing analyses, the performance of a framing detection model is tested by predicting the political stance of an article or political affiliation of a politician’s tweet or speech. However, such stance information is rarely incorporated into the development of the actual frame detection model.

We explore ways to take advantage of political stance data to improve framing analysis. The first method separately trains a stance prediction model and computes mean attention weights (MAW), which signify the reasoning behind the prediction. We use the scores of MAW to delineate important words in stance prediction. The second method is to jointly train a Transformer encoder with a contrastive learning objective for frames embedding and cross-entropy for a political stance prediction. The goal is to shift embeddings in the same framing group closer together, while increasing the distance to the language used by opposing political parties.

Our main contributions are as follows: (1) We compare the two methods proposed above to integrate stance prediction with framing analysis and investigate the effectiveness of stance prediction as a method to demonstrate the performance of a framing detection framework. (2) We show that rather than separately training stance prediction and frames embeddings, jointly training them in a multi-task learning approach better dissociates framing languages used in liberal and conservative U.S. news media.

2 Related Work

Traditionally, social scientists have developed and manually annotated a topic-specific codebook of frames (Terkildsen and Schnell, 1997; Baumgartner et al., 2008; Card et al., 2015). Computational linguists recently have applied ML techniques to

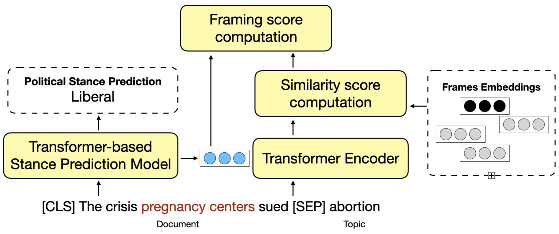

(a) Method 1: Framing Scores

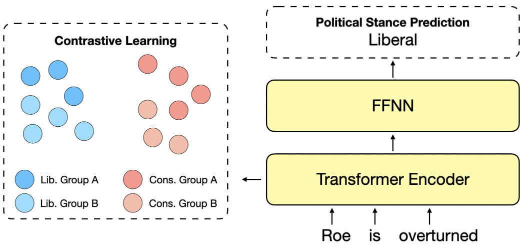

(b) Method 2: Framing Embedding

Figure 1: Two proposed frameworks. The first method is an ensemble of models to perform two separate tasks. The second method is a joint model that shares layers to learn with auxiliary tasks.

analyze frames. These works can be categorized according to their ML approaches: unsupervised, supervised, or weakly-supervised.

Many unsupervised learning approaches for frame detection depend on Latent Dirichlet Allocation (LDA) topic modeling (Blei et al., 2003) to extract candidate words of frames. However, the output of LDA is a list of keywords in each topic, *not* frame. Hence, based on the output, researchers build framing categories, i.e., frames. For instance, an open-sourced tool built by Bhatia et al. (2021) outputs the result of LDA topic modeling so that the user can label frames with the result. These topic-based words are useful guidance in framing annotations but are not appropriate data to be used for supervised framing analysis.

Second, framing detection can be defined as a supervised learning problem. Researchers collect and annotate data and train an ML model to classify frames. Field et al. (2018) constructed framing lexicons, following the Media Frame Corpus (MFC) (Card et al., 2015) annotations, and classified issue frames in Russian news articles. Akyürek et al. (2020) used the BERT (Devlin et al., 2018) to identify multilingual frames in articles about U.S. gun violence. Similar to our method, Cabot et al. (2020) applied multi-task learning to model political perspectives in news articles, political affiliations of politicians, framing, metaphor, and emotion. Our frame embedding approach differs in that the main task of our model is to embed language used in frames with contrastive learning.

Finally, there are weakly supervised models. In addition to the dictionary of frame indicators, Johnson et al. (2017b,a) used linguistic features of a text to predict framing in political tweets. Roy and Goldwasser (2020) built topic-specific lexicons by extending the lexicons in the MFC, and generalized them by creating an embedding space. Our model also creates an embedding space but uses political stance prediction as an auxiliary task and applies contrastive learning.

Based on the review of the literature, we propose two frameworks to incorporate political stance information into frame modeling. Section 3 presents our first modeling approach in which we separately train a stance prediction model and a framing encoder. Then we add feature attributions from each model to compute a final framing score. Section 4 presents the second model, which applies multi-task learning and jointly trains on political party and framing data.

3 Method 1: Framing Scores

We propose a model for frame detection that computes framing scores with feature attributions from a BERT-based stance prediction model (3.1) and phrase similarity with frame indicators (3.2). Figure 1a illustrates this approach. Section 3.1 explains how stances are calculated with MAW. Section 3.2 details how phrase similarity is computed.

3.1 Stance Rationales

We use the approach of Jayaram and Allaway (2021) to extract feature attributions, specifically mean attention weights (MAW), of their stance prediction model. The framework is a BERT-based encoder trained with an additional loss term designed to impose a prior based on human rationales. That is, the prior loss term encourages the model attributions to be similar to oracle attributions, which are important-word scores based on human annotations. After stance prediction, MAW are extracted as stance rationales. The MAW of a token j is the mean of all attention weights at index j.

3.2 Phrase Similarity

We extract candidate phrases and compute the cosine similarity between those phrases and frame indicators with Phrase-BERT (Wang et al., 2021). Phrase-BERT is a framework that fine-tunes BERT with a contrastive learning objective. We chose Phrase-BERT to embed phrases because while other BERT-based models rely on a lexical overlap to find similar phrases, Phrase-BERT is better at discovering semantically equivalent and lexically diverse phrases.

For candidate extraction, we use the implementation in EmbedRank (Bennani-Smires et al., 2018). It uses Stanford CoreNLP to identify the Part-of-Speech of each word and then generate noun phrases. We use these noun phrases as candidates of framing phrases.

Next, the candidate phrases are embedded with Phrase-BERT, and cosine similarity between a candidate and its nearest-neighbor frame indicators are computed.

The final framing score of a token is a weighted sum of MAW and the similarity score:

$$\text{MAW} \times d + SimScore \times (1 - d)$$

where d is a hyperparameter.

4 Method 2: Framing Embedding

We propose a multi-task learning framework that trains a BERT-based encoder, as shown in Figure 1b. Section 4.1 discusses how to select contrastive examples for framing embedding. Section 4.2 explains the stance prediction model.

4.1 Contrastive Learning

Similar to Phrase-BERT, we fine-tune BERT to encourage the embeddings of frame indicators in the same framing group to be close and the embeddings of frame indicators in a different group to be distant. For every frame indicator p_i, there exists a positive example p_i^+ and a negative example p_i^-. The positive examples of p are other frame indicators in the same framing group. In general, negative examples are randomly chosen from phrases or sentences that do not contain p_i. However, we specifically selected examples from framing groups that are frequently used by the opposing political party. The goal is to isolate frame indicators that are used mostly by liberal media and those that are used by conservative media. The criterion for choosing the political stance of each framing group can be found in Appendix A.

Given a triplet of vectors (p, p^+, p^-), the contrastive loss is computed as follows:

$$\mathcal{L}_c = \max(0, \epsilon - \|p - p^-\| + \|p + p^+\|)$$

where a margin ϵ is a hyperparameter.

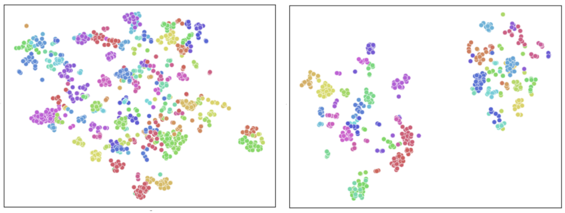

(a) Contrastive Learning (b) Framing Embedding

Figure 2: t-SNE visualiation of the embeddings of frame indicators. Figure (a) is the embeddings of BERT fine-tuned with a contrastive learning objective. Figure (b) is that of Method 2, which jointly trains with both contrastive learning and stance prediction objectives.

4.2 Stance Prediction

For stance prediction, we add a single-layer Neural Network to the BERT-based encoder. Given a frame indicator p_i, we predict a stance label $y_i \in \{0, 1\}$. We use binary cross-entropy as the loss function:

$$\mathcal{L}_s = -\frac{1}{N}\Sigma_{i=1}^{N} y_i \cdot \log(p(y_i)) + (1 - y_i) \cdot \log(1 - p(y_i))$$

where N is the number of data in a batch.

The final loss is

$$\mathcal{L} = \mathcal{L}_c + \alpha \cdot \mathcal{L}_s$$

where α is a hyperparameter.

5 Experiments

5.1 Data

We use the dataset from Roy and Goldwasser (2020). This dataset has 21,645 news articles on three politically polarized topics: abortion, immigration, and gun control. Each article is labeled *left* or *right* according to mediabiasfactcheck.com. There are also topic-specific lexicons, which were collected as in Field et al. (2018). We use these lexicons as our framing indicators.

We build a triplet dataset with the framing indicators, following the procedure explained in Section 4.1. The dataset has 7,366 triplets. Unlike Roy and Goldwasser (2020), we do not create separate embeddings for each topic; we embed framing indicators from three topics into one embedding space.

5.2 Results

We evaluate the performance of our first model by highlighting words in a document according to their scores. Table 1 shows highlights from an article on abortion. The darker the highlight is, the higher

MAW	Abortion rights advocates even say that the legislation could lead to the end of private insurance coverage for abortion . As I reported Susan Cohen the director of Governmental Affairs for the pro-abortion-rights Guttmacher Foundation argued in a policy brief this fall that the Smith Bill would go into uncharted.
SimScore	Abortion rights advocates even say that the legislation could lead to the end of private insurance coverage for abortion. As I reported Susan Cohen the director of Governmental Affairs for the pro-abortion-rights Guttmacher Foundation argued in a policy brief this fall that the Smith Bill would go into uncharted .
Framing	Abortion rights advocates even say that the legislation could lead to the end of private insurance coverage for abortion. As I reported Susan Cohen the director of Governmental Affairs for the pro-abortion-rights Guttmacher Foundation argued in a policy brief this fall that the Smith Bill would go into uncharted .

Table 1: The visualization of mean attention weights (MAW), cosine similarity scores (SimScore), and the final framing scores. The example document is an article with the topic Abortion.

Phrase	**Nearest Neighbors**
Protecting the preborn	baby's life, child's life, kill the child, child protection, kill the baby, unborn life, child killing, abort the baby, rip the baby, protect life
Prevent firearm violence	prevent gun violence, curb gun violence, gun violence restraining, end gun violence, stop gun violence, violence restraining order, reduce gun death, domestic violence restraining, gun violence research, violence restraining
Illegal immigrants are criminals	deport illegal immigrant, previously deported illegal, deport illegal, amnesty to illegal, deport undocumented, terrorist organization, deportation of illegal, domestic terrorism, suspected terrorist, terrorism related

Table 2: Top 10 nearest neighbors of phrases that were not in the dataset.

the score is. As shown in Table 1, important tokens based on MAW and SimScore are distinguishable. However, MAW scores were mostly proportionate across all tokens, and thus the effects of MAW on final framing scores were statistically insignificant.

For our second model, we used t-SNE to visualize the embeddings of frame indicators. Figure 2a shows the embedding space of BERT fine-tuned with our contrastive learning objective. Figure 2b shows the embedding space of the Method 2 model, which added a stance prediction loss. The *left* and *right* separation of framing groups is evident in Figure 2b. This result suggests that using stance prediction as an auxiliary task improves the embedding of framing.

Next, we evaluated the performance of the second model in embedding previously unseen phrases. Table 2 shows the top 10 nearest neighbors of those phrases. The phrases include vocabularies that were not present in the dataset. For instance, the word "preborn" in the phrase "protecting the preborn" was not present in the framing indicators. Still, the model was able to assign its embedding close to framing indicators that do not have lexical overlaps but are semantically similar to frames used by conservative labeled articles. Yet the nearest neighbors were restricted to existing framing groups; that is, the model could not extend the assignment of phrases to unobserved framing.

6 Conclusion

In this work, we proposed and compared two frameworks that incorporate stance prediction to framing detection and have shown initial results that jointly learning the two tasks is the strongest model. We plan to extend this work with quantified experiments to discover frames in unlabeled data.

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*.

Frank R Baumgartner, Suzanna L De Boef, and Amber E Boydstun. 2008. *The decline of the death penalty and the discovery of innocence*. Cambridge University Press.

Kamil Bennani-Smires, Claudiu Musat, Andreea Hossmann, Michael Baeriswyl, and Martin Jaggi. 2018. Simple unsupervised keyphrase extraction using sentence embeddings. *arXiv preprint arXiv:1801.04470*.

Vibhu Bhatia, Vidya Prasad Akavoor, Sejin Paik, Lei Guo, Mona Jalal, Alyssa Smith, David Assefa Tofu, Edward Edberg Halim, Yimeng Sun, Margrit Betke, et al. 2021. Openframing: Open-sourced tool for computational framing analysis of multilingual data. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 242–250.

David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022.

Pere-Lluís Huguet Cabot, Verna Dankers, David Abadi, Agneta Fischer, and Ekaterina Shutova. 2020. The pragmatics behind politics: Modelling metaphor, framing and emotion in political discourse. *ACL Anthology*.

Dallas Card, Amber Boydstun, 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.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Robert M Entman. 1993. Framing: Towards clarification of a fractured paradigm. *McQuail's reader in mass communication theory*, pages 390–397.

Robert M Entman and Andrew Rojecki. 1993. Freezing out the public: Elite and media framing of the us anti-nuclear movement.

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. *arXiv preprint arXiv:1808.09386*.

Lei Guo, Chris J Vargo, Zixuan Pan, Weicong Ding, and Prakash Ishwar. 2016. Big social data analytics in journalism and mass communication: Comparing dictionary-based text analysis and unsupervised topic modeling. *Journalism & Mass Communication Quarterly*, 93(2):332–359.

Sahil Jayaram and Emily Allaway. 2021. Human rationales as attribution priors for explainable stance detection. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5540–5554.

Kristen Johnson and Dan Goldwasser. 2016. "all i know about politics is what i read in twitter": Weakly supervised models for extracting politicians' stances from twitter. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2966–2977.

Kristen Johnson, Di Jin, and Dan Goldwasser. 2017a. Leveraging behavioral and social information for weakly supervised collective classification of political discourse on twitter. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 741–752.

Kristen Johnson, I-Ta Lee, and Dan Goldwasser. 2017b. Ideological phrase indicators for classification of political discourse framing on twitter. In *Proceedings of the Second Workshop on NLP and Computational Social Science*, pages 90–99.

Matthew C Nisbet, Dominique Brossard, and Adrianne Kroepsch. 2003. Framing science: The stem cell controversy in an age of press/politics. *Harvard International Journal of Press/Politics*, 8(2):36–70.

Shamik Roy and Dan Goldwasser. 2020. Weakly supervised learning of nuanced frames for analyzing polarization in news media. *arXiv preprint arXiv:2009.09609*.

Nayda Terkildsen and Frauke Schnell. 1997. How media frames move public opinion: An analysis of the women's movement. *Political research quarterly*, 50(4):879–900.

Shufan Wang, Laure Thompson, and Mohit Iyyer. 2021. Phrase-bert: Improved phrase embeddings from bert with an application to corpus exploration. *arXiv preprint arXiv:2109.06304*.

Philipp Wicke and Marianna M Bolognesi. 2020. Framing covid-19: How we conceptualize and discuss the pandemic on twitter. *PloS one*, 15(9):e0240010.

Framing group	Lib.	Cons.
Pro-Life	0.563	0.437
Sanctity of Life	0.554	0.446
Pro-Choice	0.791	0.209
Right of Human Life	0.361	0.639
Life Protection	0.429	0.571
Abort. prov. economy	0.096	0.904
Reproduction Right	0.829	0.171
Sale of Fetal Tissue	0.57	0.43
Late Term Abortion	0.531	0.469
Abortion Funding	0.621	0.379
Hobby Lobby	0.714	0.286
Anti-Abortion	0.723	0.277
Health Care	0.818	0.182
Women freedom	0.786	0.214
Roe V. Wade	0.615	0.385
Birth Control	0.689	0.311
Planned Parenthood	0.426	0.574
Sexual Assault Vict.	0.649	0.351
Pregnancy Centers	0.295	0.705
Stem Cell Research	0.727	0.273

Table 3: Topic: Abortion. Proportion of liberal and conservative articles that mention frame indicators in each category.

A Stance Labelling of Framing Groups

Tables 3, 4, and 5 show the usage of framing in liberal and conservative articles of the Roy and Goldwasser (2020) dataset. If frame indicators in a group were mentioned more in liberal media, that group is labeled as "left." Similarly, a group, which is more used in conservative media, is labeled as "right."

B Reproducability

Machine Used We used a Nvidia Quatro RTX 5000, 16 GB memory GPU in a machine with Intel(R) Xeon(R) Silver 4214 CPU @ 2.20GHz.

Libraries Used For all models, PyTorch was used for implementation.

Implementation Details

- **Method 1: Stance Rationales** The hyper-parameter λ is 49152. The model is optimized using Adam for 20 epochs with a batch size of 32 and a fixed learning rate of 10^{-5}. The maximum sequence length of 250 for arguments and 10 for topics. It

Framing group	Lib.	Cons.
Gun Buyback Program	0.781	0.219
Terrorist Attack	0.441	0.559
Gun Con. to Restr. Viol.	0.628	0.372
White Identity	0.761	0.239
Gun Research	0.843	0.157
Mental Health	0.622	0.378
Gun Show Loophole	0.55	0.45
Gun Business Industry	0.513	0.487
Second Amendment	0.358	0.642
Assault Weapon	0.466	0.534
Person of Color Identity	0.686	0.314
Conc. Carry Recip. Act	0.401	0.599
Gun Homicide	0.692	0.308
Ban on Handgun	0.471	0.529
Right to Self-Defense	0.275	0.725
School Safety	0.644	0.356
Background Check	0.496	0.504
Stop Gun Crime	0.535	0.465
Illegal Gun	0.446	0.554

Table 4: Topic: Gun Control. Proportion of liberal and conservative articles that mention frame indicators in each category.

Framing group	Lib.	Cons.
Terrorism	0.417	0.583
Born identity	0.156	0.844
Human Right	0.728	8 0.272
Wage Economy	0.653	0.347
DACA	0.455	0.545
Detention	0.696	0.304
Deport.: In General	0.58	0.42
Salary Stagnation	0.137	0.863
Rac. and Xen.	0.667	0.333
Border Protection	0.411	0.589
Cheap Labor Availability	0.032	0.968
Wealth Gap	0.106	0.894
Refugee	0.482	0.518
Taxpayer Money	0.461	0.539
Amnesty	0.034	0.966
Racial Identity	0.837	0.163
Deport.: Ill. Imm.	0.184	0.816
Birth Cit. & 14th Amen.	0.435	0.565
Merit Based Imm.	0.1	0.9
Dream Act	0.342	0.658
Asylum	0.519	0.481
Family Sep. Policy	0.761	0.239

Table 5: Topic: Immigration. Proportion of liberal and conservative articles that mention frame indicators in each category.

uses `bert-base-uncased` from Huggingface.[1]

- **Method 1: Phrase-BERT** The model is optimized using Adam for 1 epoch with a batch size of 16 and a learning rate of $2e-5$. The initial 10% of training steps are used as warm-up steps. It uses `bert-base-nli-stsb-mean-tokens` from Huggingface.

- **Method 1: Framing Score** The hyperparemeter d is set as $d \in \{0, 0.1, 0.2, 0.3, 0.5, 0.7, 0.8, 1\}$.

- **Method 2: Framing Embedding** The hyperparmaeter α is set as 0.2. The model is optimized using Adam for 10 epoch with a batch size of 16 and a learning rate of $2e-5$. The initial 10% of training steps are used as warm-up steps.

C Ethical Considerations

As mentioned in Section 5, our methods have limitations, and we caution not to deploy our models for not the purpose intended.

[1] https://huggingface.co/transformers