# CLoSE: Contrastive Learning of Subframe Embeddings for Political Bias Classification of News Media

**Michelle YoungJin Kim** and **Kristen Marie Johnson**
Computer Science and Engineering, Michigan State University, USA
{kimmic16,kristenj}@msu.edu

## Abstract

Framing is a political strategy in which journalists and politicians emphasize certain aspects of a societal issue in order to influence and sway public opinion. Frameworks for detecting framing in news articles or social media posts are critical in understanding the spread of biased information in our society. In this paper, we propose CLoSE, a multi-task BERT-based model which uses contrastive learning to embed indicators of frames from news articles in order to predict political bias. We evaluate the performance of our proposed model on subframes and political bias classification tasks. We also demonstrate the model's classification accuracy on zero-shot and few-shot learning tasks, providing a promising avenue for framing detection in unlabeled data.

## 1 Introduction

News media coverage shapes our attitudes, emotions, and decisions toward public issues (Iyengar, 1994; Pan and Kosicki, 1993; Jensen et al., 2014). Research shows that people's perceptions of news can be manipulated by presenting the same story with different expressions. For example, participants of a study found a terrorist attack caused by "al-Qaeda and associated radical Islamic groups" considerably more threatening than a terrorist attack by "domestic rebel separatist groups," which is an equivalent paraphrase (Kapuściński and Richards, 2016). Hence, studies on the influences of different presentations of issues, or the effects of *framing*, play an essential role in understanding political discourse.

Framing refers to emphasizing desired aspects of an issue to promote and amplify a particular perspective (Entman, 1993). By selecting and thus elevating the salience of a specific angle of a topic, media sources can present the topic through their choice of frames to induce particular attributes and judgments among the public. Framing is widely researched on various topics, from its effects on public opinion on political issues such as the U.S. anti-nuclear war movement (Entman and Rojecki, 1993) and stem cell research (Nisbet et al., 2003) to the economic impact of framing (Liu and Pennington-Gray, 2015; Van Dalen et al., 2017).

In this work, we study framing detection for three politically polarized issues in the U.S. news media: abortion, gun control, and immigration. We focus on framing in the news discourse to understand a discrepancy in media consumption and its influence on political bias polarization. With the increase of news media outlets and social media platforms, the public is overwhelmed with a flood of information. Unfortunately, the increase of news sources does not yield the sharing of information across political views, often developing into biased echo chambers on social media platforms. In fact, there are stark differences in the social media and news sources that liberals and conservatives use and trust. Fox News is the primary news source for nearly half of conservatives, while NPR, MSNBC, and the New York Times are the most trusted news sources for liberals (Mitchell et al., 2014).

Machine learning techniques have been applied to detect and analyze political frames (Card et al., 2015; Guo et al., 2016; Johnson et al., 2017a; Bhatia et al., 2021). In such framing analyses, the performance of a framing detection model is tested by predicting the political bias of an article or the political affiliation or stance of a politician's tweet or speech. However, we suggest incorporating the political bias information into the development of the actual frame detection model. We focus on general liberal and conservative bias rather than the specific stances politicians may take on issues.

In this paper, we propose CLoSE, a framing embedding model that jointly learns to predict political bias. CLoSE fine-tunes BERT variants with a contrastive learning objective to generate (sub)frame embeddings for a given input sentence. Then we



*Proceedings of the 29th International Conference on Computational Linguistics*, pages 2780–2793
October 12–17, 2022.

add a prediction loss to classify the political bias (liberal or conservative) of the embedded text.

For the embedding task, we use contrastive learning to place embeddings of the same *subframe* closer together. Subframes are fine-grained subcategorizations of the general political frames of Boydstun et al. (2014). The topic-specific lexicons of subframes are the *subframe indicators* used to identify specific framing language within a chosen topic. We use the subframe indicators to identify texts with frames and construct the Framing Triplet Dataset. This dataset, built explicitly for a contrastive objective, consists of a triplet of an anchor sentence, its positive sample, and its negative sample, and the political bias label of the anchor sentence. The contrastive learning objective reduces the distance between the anchor sentence and its positive sample, which belongs to the same subframe, while increasing the distance to its negative sample that belongs to a different subframe.

CLoSE outperforms the baseline models in both the subframe and political bias classification tasks. The results also show that the contexts of subframes improve the performance of the political bias classifier. Further, our pre-trained model accomplishes superior performance in zero-shot and few-shot settings. Finally, we design a topic modeling method for the subframe embeddings that clusters nearby embeddings, extracts statistically significant words with class-based TF-IDF (Grootendorst, 2022), and merges clusters with high overlapping words.

To summarize, our main contributions are as follows: (1) We collect and release the Triplet Framing Dataset[1], a triplet of sentences that include subframe indicators. (2) We propose CLoSE, contrastive learning of subframe embeddings model that jointly generates embeddings and predicts political bias of framed texts. The experiments demonstrate our method's performance on subframe and political bias classification tasks and investigate the effectiveness of political bias information on subframe detection and vice versa. (3) The experiments show that our pre-trained model performs competitively on zero-shot political bias classification and few-shot subframe classification tasks. Namely, the pre-trained model can predict the political bias of articles with previously unseen topics and predict subframe groups with a limited quantity of labeled data. This greatly reduces the cost of data annotation for framing and bias tasks.

[1]Code and data available at https://github.com/MSU-NLP-CSS/CLoSE_framing.

## 2 Related Work

Framing is a powerful political strategy that is used to influence public opinion. Hence identifying what and how frames are used is a critical task in political communications. Framing in news media and social networks has been studied to analyze political polarization (Johnson and Goldwasser, 2016; Tsur et al., 2015; Tourni et al., 2021). However, annotating data for framing is challenging due to the nuanced language of frames across political issues.

To overcome this challenge, computational social scientists follow a topic-specific codebook to manually annotate documents with frames (Terkildsen and Schnell, 1997; Baumgartner et al., 2008; Card et al., 2015). The most commonly used codebook is the Policy Frames Codebook (Boydstun et al., 2014), which proposes a general coding scheme for fifteen high-level frames across policy issues. Based on this codebook, Card et al. (2015) collected and released the Media Frames Corpus (MFC). The MFC contains more than 20,000 news articles on three policy issues: immigration, smoking, and same-sex marriage.

Following the release of MFC, the first large-scale open-source frames dataset, many researchers studied and analyzed frames by leveraging the MFC. Johnson et al. (2017b) use political tweets to extract phrases that frequently appear in each of the frames and propose a framing detection model that uses both linguistic features and the extracted ideological phrases. Field et al. (2018) extend the U.S. policy frames to Russian policy frames and analyze 13 years of Russian news articles. They derive the lexicons of each frame from the MFC and translate them to Russian to generate Russian framing lexicons. Huguet Cabot et al. (2020) propose a multi-task model that learns metaphor, framing, and emotion in political discourse and uses the MFC to predict frames. Although our method also jointly models political bias and framing, our approach differs in that the main task of our model is to embed language used in subframes with contrastive learning.

Additional research has narrowed down the general policy frames to suggest issue-specific frames. Johnson et al. (2017a) add Twitter-specific frames to the policy frames and annotate and analyze the tweets of U.S. politicians. Roy and Goldwasser (2020) build topic-specific lexicons, which they de-

fine as *subframe indicators*, by using an embedding model to generalize the MFC lexicon for analyzing media bias. We utilize their subframe indicators to create our Framing Triplet Dataset. Our proposed model also includes an embedding model that generalizes text containing subframe indicators. However, we integrate political bias information into our model directly by adding bias classification as an auxiliary task.

Other approaches for detecting frames are based on topic modeling algorithms. Latent topics of a given corpus are extracted with a topic modeling algorithm, often Latent Dirichlet Allocation (LDA) (Blei et al., 2003). LDA yields statistically significant words of each topic, which are used as candidate indicators for defining frames. Bhatia et al. (2021) provide an open-sourced tool for analyzing frames in multilingual texts. Given text inputs, their web-based system sends the LDA topic modeling output to a user's email so that the user can decide and label frames on the given topics. However, we note that the output of LDA is a list of keywords in each topic, *not* frame. The topic-based words are helpful guidance in framing annotations and exploratory analysis but are not appropriate data to be used for supervised framing analysis.

## 3 Framing Triplet Dataset

We introduce the Framing Triplet Dataset, which contains 25,627 news articles on three politically polarized topics: abortion, gun control, and immigration. We extend the hyper-partisan news dataset (Kiesel et al., 2019) and the dataset by Roy and Goldwasser (2020), which consists of news articles on the three topics until 2019. Open-source crawlers `pygooglenew` and `news-please` were used to collect recent articles from 2020 to 2022. We query the Google News RSS feed with keywords of each topic and crawl the articles. The keywords are "abortion" for the abortion data, "gun" for the gun control data, and "immigrant" and "immigration" for the immigration data. Then we assign the political bias label of each article according to Media Bias Fact Check [2], the largest crowdsourced media bias resource on more than 4,500 media sources and journalists. The Media Bias Fact Check categorizes each media source into one of the following five biases: "left bias," "left-center bias," "least bias," "right-center bias," and "right bias." For our study, given the partisan divide of current U.S. politics, we only consider articles from "left bias" and "right bias" media. In total, our dataset includes 12,911 left-biased articles and 12,761 right-biased articles. The detailed statistics of the dataset can be found in Table 1.

After labeling the political bias of the collected articles, we extract headlines and sentences in articles that contain *subframe indicators*. The subframe indicators, suggested by Roy and Goldwasser (2020), are the issue-specific subclassifications of the policy frames by Boydstun et al. (2014). These subframe indicators are collected via the following three steps. First, the top-250 unigrams for each of the 15 policy frames are gathered based on their Pointwise Mutual Information (PMI) scores (Church and Hanks, 1990). Then, each paragraph in the article is annotated with the policy frames based on the number of unigram matches. Finally, repeating phrases in the annotated paragraphs are grouped to form subframes, which represent topic-specific lexicons. They defined 20 subframes for the topic of abortion, 22 subframes for the topic of immigration, and 19 subframes for the topic of gun control. The list of subframes can be found in Appendix A. We refer to Roy and Goldwasser (2020) for the full list of subframe indicators.

Finally, we construct triplets with the extracted sentences. Given an anchor sentence $s_a$, we define its positive sample $s_p$ as a sentence with subframes from the same subframe group and its negative sample $s_n$ as a sentence with subframes from a different subframe group. For example, the sentence "The first backlash to *the Roe decision* came primarily from groups representing U.S. Catholics." contains a subframe indicator "Roe decision" that belongs to the subframe "Roe v. Wade." Its positive sample ($s_p$) should be a sentence with a subframe indicator of that same subframe, "Roe v. Wade." Its negative example ($s_n$) will be a sentence with a subframe indicator from a different, random subframe group such as "Birth Control." The triplet is formed as $(s_a, s_p, s_n)$. An example of the triplet data is shown in Figure 2.

### 3.1 Data Analysis

We analyze the usage of subframes in the Framing Triplet Dataset by extracting the top-3 subframes across time and issues. The results for the abortion subdataset are shown in Table 2. (See Appendix B for the results on the gun control and immigration subdatasets.) Across the topics, we notice the trend

[2] `mediabiasfactcheck.com`

| | News(#) | Left(#) | Right(#) | Sent.(#) | Left S.(#) | Right S.(#) | Time Span |
|---|---|---|---|---|---|---|---|
| **Abortion** | 8,061 | 4,518 | 3,543 | 10,725 | 8,032 | 2,693 | 1984-2022 |
| **Gun control** | 9,476 | 4,238 | 5,238 | 8,138 | 3,918 | 4,220 | 2000-2022 |
| **Immigration** | 8,090 | 4,155 | 3,935 | 13,269 | 8,228 | 5,041 | 1996-2022 |

Table 1: Statistics of the Framing Triplet Dataset. For each topic, news articles are collected, their biases are labeled according to Media Bias Fact Check, and sentences containing subframe indicators are extracted. *Left S.* are sentences extracted from left-biased news articles. *Right S.* are extracted from right-biased news articles.

| Years | | Liberal | Conservative |
|---|---|---|---|
| 2020-2022 | Top-1 | Roe v. Wade | Roe v. Wade |
| | Top-2 | Abortion Funding | Abortion Funding |
| | Top-3 | Birth Control | Birth Control |
| 2018-2020 | Top-1 | Roe v. Wade | Roe v. Wade |
| | Top-2 | Birth Control | Pregnancy Centers |
| | Top-3 | Abortion Funding | Abort. Prov. Economy |
| 2016-2018 | Top-1 | Birth Control | Pregnancy Centers |
| | Top-2 | Health Care | Abort. Prov. Economy |
| | Top-3 | Roe v. Wade | Abortion Funding |
| 2010-2016 | Top-1 | Abortion Funding | Birth Control |
| | Top-2 | Birth Control | Abort. Prov. Economy |
| | Top-3 | ⋆ | Planned Parenthood |
| -2010 | Top-1 | Birth Control | Roe v. Wade |
| | Top-2 | Roe v. Wade | Right of Human Life |
| | Top-3 | ⋆ | Abortion Funding |

⋆ Potentially offensive or triggering language has been omitted.

Table 2: Top-3 Subframe Indicators of the Abortion Triplet Dataset. Each row indicates the most frequently used subframe indicators for liberal and conservative biased media within the specified time frame.

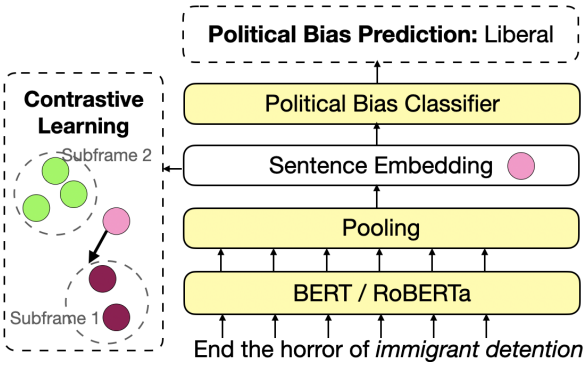


Figure 1: Architecture of CLoSE. Our model consists of a BERT- or RoBERTa-based encoder whose output is fed through a pooling layer to generate sentence embeddings via a contrastive learning objective. These embeddings are then used in the final political bias classification task. This example shows an input sentence and its corresponding embedding (pink circle), which is most similar to other embedded sentences of the Subframe 1's group. The final prediction here is liberal.

that the most used subframes begin to overlap as time passes. Before 2018, the top-3 subframes used by the liberal and conservative media significantly differed. Namely, the subframe indicators that the media focused on to frame people's opinions on abortion are different based on the media's political bias. However, the most used subframes of the left and right news are intersecting more recently. From 2018 to 2020, the most used subframe of both liberal and conservative media was "Roe v. Wade." The top-3 subframes from 2020 to 2022 of both views are also identical. Similarly, from 2020 to 2022, the top-3 subframes for the gun control subdataset, and the top-2 subframes for the immigration subdataset, of both liberal and conservative media are identical. These results show that in recent years, news media with opposing political biases cover common issues, sometimes with similar framing language. Despite this similar coverage, political polarization in the U.S. remains dominant (Doherty et al., 2021). Hence, we need a methodology that captures not only which issues the media chooses to spotlight for coverage but also how they frame these issues for public perception. By jointly learning to predict political bias from the sentence embeddings, our proposed method aims to capture the different contexts of the opposing media.

## 4 Methodology

We propose CLoSE: a multi-task learning model that jointly learns to embed sentence framing language and predict political bias. As shown in Figure 1, the model consists of a BERT-based or RoBERTa-based encoder, followed by a pooling layer that generates a sentence embedding, which feeds into a classifier for political bias prediction. We adopt a contrastive learning objective so that a sentence with subframe indicators close to other sentences of the same subframe and far from those of different subframes will be similarly reflected in the embedding space.

**Embedding via pooling.** Given a sentence $s$ of tokens $\{t_1, \cdots, t_n\}$, we embed the sentence using BERT or RoBERTa to get a sequence of token embeddings $\{e(t_1), \cdots, e(t_n)\}$. Then a pooling operation is applied to the token embeddings to

derive a fixed size sentence embedding $e(s)$. There are three possible pooling strategies: (1) taking the output of the CLS token, (2) computing the mean of all output vectors, and (3) computing a max-over-time of the output vectors. CLoSE employs the default pooling operation of mean strategy, which has shown the best experimental results on semantic textual similarity and natural language inference tasks (Reimers and Gurevych, 2019). Thus, the final sentence embedding is $e(s) = \frac{1}{n}\sum_{i=1}^{n} e(t_i)$.

**Learn subframes with a contrastive learning objective.** We apply a contrastive learning objective to capture the framing embedding space. To do so, we fine-tune the BERT-based (or RoBERTa-based) encoder to encourage embeddings that will fall within the same subframe group to be closer and the embeddings within a different subframe group to be more distant from a given anchor sentence's embedding. Every anchor sentence $s_a$ has a positive sentence $s_p$ and a negative sentence $s_n$, corresponding to the triplets within the Framing Triplet Dataset. For example, as shown in Figure 2, the anchor sentence has a subframe indicator "Roe decision" (shown in bold), which belongs to the subframe group "Roe v. Wade." While its positive sentence has subframe indicators "overturns Roe" and "Roe v. Wade" that belong to the same subframe group, the negative sentence has a subframe indicator "March for Life" that belongs to a different subframe group "Pro-Life."

Given a triplet of sentences $(s_a, s_p, s_n)$, the contrastive learning loss is computed as follows:

$$\mathcal{L}_c = \max(0, \epsilon - \|e(s_a) - e(s_n)\| + \|e(s_a) - e(s_p)\|)$$

where the margin $\epsilon$ is a hyperparameter, and $\|\cdot\|$ is the Euclidean distance. $e(s_a), e(s_p)$, and $e(s_n)$ are an embedding of the anchor sentence $s_a$, its positive sentence $s_p$, and its negative sentence $s_n$, respectively. We set $\epsilon$ as 1 for the experiments.

**Predict political bias.** The proposed model adds a binary classifier to the output of the BERT-based encoder to predict the political bias of an input (anchor) sentence. The sentence embedding of this anchor sentence is given as an input to the classifier, which predicts whether the sentence is from a liberal (left) or conservative (right) news source. Given an anchor sentence $s_i$, we predict a stance label $y_i \in \{0, 1\}$. We use binary cross-entropy as the loss function:

$$\mathcal{L}_{\text{BCE}} = -\frac{1}{N}\sum_{i=1}^{N} y_i \cdot \log(p(y_i)) + (1 - y_i) \cdot \log(1 - p(y_i))$$

where $N$ is the number of data in a batch. In our experimental studies, the classifier is a single-layer feedforward neural network with Dropout (Srivastava et al., 2014) and ReLU activation (Nair and Hinton, 2010).

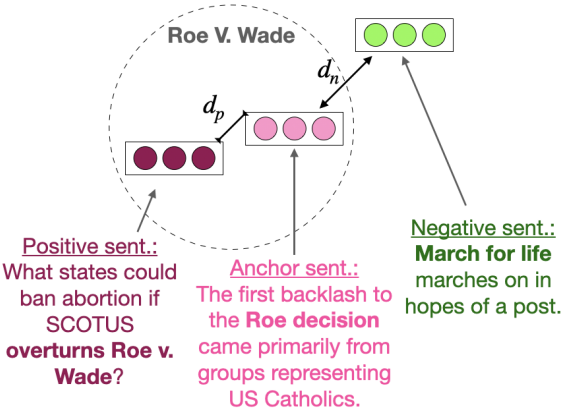


Figure 2: Visualization of the Contrastive Learning Objective. The contrastive learning objective pushes the anchor sentence closer to its positive sentence (decreasing distance to positive example $d_p$) and farther from its negative sentence (increasing distance from $d_n$).

**Jointly learn framing and political bias.** Finally, we combine the learned losses in order to build an embedding space that can group subframes as well as separate the contexts of opposing political stances, i.e., distinguish between liberal and conservative biased news. The final loss that computes the weighted sum of the contrastive learning loss $L_c$ and the classification loss $L_{BCE}$ is:

$$\mathcal{L} = (1 - \alpha) \cdot \mathcal{L}_c + \alpha \cdot \mathcal{L}_{\text{BCE}} \quad (1)$$

where $\alpha$ is a hyperparameter. The default value of $\alpha$ is 0.5.

## 5 Experimental Setup

We evaluate the performance of our proposed method on the tasks of (1) political bias classification and (2) subframe detection. For the first task, a model makes a binary prediction of whether a given text has a liberal or conservative perspective. We compare the classification performance of our proposed model to the baseline models. Further,

an ablation study on the classification objective is performed, and the classification performance of the pre-trained models on unseen topic data is studied. As for the second task, we test the performance of subframe classification and analyze the effect of the bias classification task on the subframe classification. Finally, we propose a novel topic modeling method for the output embeddings and show experimental results.

### 5.1 Training Details

We leverage the pre-trained BERT and RoBERTa models as the encoder for our proposed model. We use the following pre-trained models from Hugging Face: `bert-base-cased`, `bert-base-uncased`, and `roberta-base`. They are trained on English texts with 12 layers, 768 hidden units, and 12 attention heads. We use Adam (Kingma and Ba, 2015) optimizer and set learning rates to $\{1e-5, 2e-5\}$, an epoch to 5, and a dropout rate to 0.5. For all models, PyTorch was used for implementation. All experiments are conducted on an Nvidia Quatro RTX 5000, 16 GB memory GPU in a machine with Intel(R) Xeon(R) Silver 4214 CPU @ 2.20GHz.

### 5.2 Baseline Models

We compare the performance on the political bias classification task with the following baselines.

**$\text{BERT}_{\text{cased}}$**: A pre-trained BERT with a single-layer feedforward neural network. The BERT encoder is pre-trained with case-sensitive English texts.

**$\text{BERT}_{\text{uncased}}$**: A pre-trained BERT with a single-layer feedforward neural network. The BERT encoder is pre-trained with case-insensitive English texts.

**RoBERTa**: A pre-trained RoBERTa model with a single-layer feedforward neural network.

## 6 Results & Discussion

### 6.1 Political Bias Classification

We compare the performance of CLoSE against the baseline models described in Section 5.2 on the task of political bias classification. The results are reported in Table 3. Our modeling approach outperforms the baseline models for both gun control and immigration bias classification and showed competitive results for abortion classification. This indicates that jointly modeling the embeddings with a contrastive learning objective improves bias classification.

| | **Abortion** | **Gun** | **Immig.** |
|---|---|---|---|
| $\text{BERT}_{\text{cased}}$ | 0.790 | 0.698 | 0.801 |
| $\text{BERT}_{\text{uncased}}$ | 0.788 | 0.722 | 0.824 |
| RoBERTa | **0.831** | 0.720 | 0.827 |
| $\text{CLoSE}_{\text{BERT-cased}}$ | 0.821 | 0.705 | **0.841** |
| $\text{CLoSE}_{\text{BERT-uncased}}$ | 0.803 | 0.716 | 0.818 |
| $\text{CLoSE}_{\text{RoBERTa}}$ | 0.758 | **0.724** | 0.829 |

Table 3: $F_1$ Scores for Political Bias Classification. Scores in bold indicate the best performing model for that topic.

| | **Abortion** | **Gun** | **Immig.** |
|---|---|---|---|
| $\text{CLoSE}_{\text{Abort.}}$ | — | 0.722 | 0.738 |
| $\text{CLoSE}_{\text{Gun}}$ | 0.749 | — | 0.735 |
| $\text{CLoSE}_{\text{Immig.}}$ | 0.747 | 0.686 | — |

Table 4: $F_1$ Scores for Zero-Shot Political Bias Classification. Cells with — indicate the topic subdataset the model was trained on.

| $\alpha =$ | 0 | 0.5 | 1 |
|---|---|---|---|
| Abortion | 0.297 | 0.821 | 0.787 |
| Gun Control | 0.328 | 0.705 | 0.514 |
| Immigration | 0.184 | 0.841 | 0.734 |

Table 5: Ablation Study. We compute $F_1$ scores of the political bias classification, dependent on the values of $\alpha$. For $\alpha = 0$, only the contrastive learning objective is used. For $\alpha = 1$, only the political bias classification loss is calculated. For brevity, other $\alpha$ values are omitted but reflect similar patterns.

Interestingly, our pre-trained model also performs well in a zero-shot learning setting. That is, our model can predict the political bias of articles from unseen topics. As shown in Table 4, the model $\text{CLoSE}_{\text{Abort.}}$ trained on only the abortion subdataset of the Triplet Dataset is able to predict the political stance of articles from the gun control and immigration subdatasets with $F_1$ scores above 0.7. Furthermore, the $F_1$ score for predicting bias of the gun control topic is 0.722. This score is extremely close to the best classification score on the gun subdataset by $\text{CLoSE}_{\text{RoBERTa}}$, as shown in Table 3. The performance of our modeling approach in this zero-shot learning setting indicates the model's ability to learn and transfer latent framing indicators for political bias detection across topics.

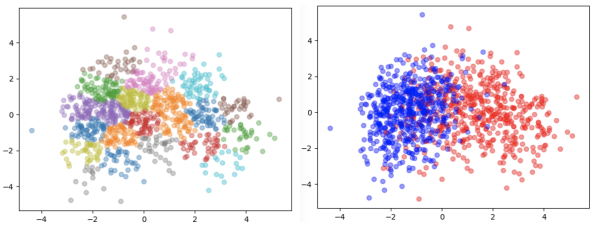

(a) Subframes (b) Political Bias

Figure 3: Visualization of the Embeddings for Immigration. Subfigure (a) shows the clustering of subframe groups. Subfigure (b) is labeled based on the political bias of each article.

| | **Abortion** | **Gun** | **Immig.** |
|---|---|---|---|
| $\text{BERT}_{\text{cased}}$ | 0.879 | 0.761 | 0.734 |
| $\text{CLoSE}_{\text{BERT-cased}}$ | 0.848 | 0.853 | 0.905 |

Table 6: $F_1$ Scores of the Subframe Classification Task.

### 6.2 Ablation Study

To further verify the usefulness of our joint learning objective, we compute $F_1$ scores of the political bias classification dependent on the values of the $\alpha$ hyperparameter of Equation 1. For $\alpha = 0.5$, we compute the sum of the contrastive objective and classification loss. When $\alpha = 0$, the loss function becomes the contrastive learning objective. On the other hand, when $\alpha = 1$, only the political bias classification loss is considered. The results are shown in Table 5.

We notice that when only the contrastive loss is applied, the model performs poorly on the political bias classification task. As expected, the performance improves when only the political bias classification loss is used (i.e., when $\alpha$ = 1). The interesting observation is that when we use *both* the contrastive loss and the bias classification loss, as we have proposed in this work, the $F_1$ scores are the highest. These results show that incorporating subframe indicators indeed improves political bias classification.

### 6.3 Subframe Classification

We evaluate whether our proposed model can correctly predict the subframe group of a given sentence. The results are shown in Table 6. We add a classifier to the pre-trained BERT-based encoder and predict the subframe label of a given input. We use $\text{CLoSE}_{\text{BERT-cased}}$, which demonstrated competitive performance across the data in the classification task, as our pre-trained encoder and compare it to the baseline model $\text{BERT}_{\text{cased}}$. For both

| | | $\text{BERT}_{\text{cased}}$ | | $\text{CLoSE}_{\text{BERT-cased}}$ | |
|---|---|---|---|---|---|
| | | *Acc.* | $F_1$ | *Acc.* | $F_1$ |
| **Abort.** | 50 | 0.232 | 0.232 | 0.300 | 0.300 |
| | 100 | 0.454 | 0.454 | 0.641 | 0.642 |
| **Gun** | 50 | 0.265 | 0.270 | 0.485 | 0.486 |
| | 100 | 0.666 | 0.679 | 0.756 | 0.759 |

Table 7: Accuracy and $F_1$ Scores of Subframe Classification in Few-shot Learning Setting. CLoSE, pre-trained on immigration data, is fine-tuned on a limited number (50 or 100 triplets) of abortion or gun control data.

| | Original Data | $\text{CLoSE}_{\text{BERT-cased}}$ |
|---|---|---|
| **Abortion** | 0.413 (0.294) | 0.347 (0.225) |
| **Gun** | 0.303 (0.189) | 0.198 (0.182) |
| **Immig.** | 0.490 (0.297) | 0.383 (0.281) |

Table 8: Average Differences of Subframe Usage Between Political Ideologies. Standard deviations are shown in parentheses.

gun control and immigration subdatasets, our pre-trained model significantly outperforms the baseline model. For the abortion data, the $F_1$ score of the baseline model is higher; yet the difference in scores was not statistically significant.

Table 7 shows that our model outperforms the baseline model even in few-shot learning. For this experiment, $\text{CLoSE}_{\text{BERT-cased}}$, pre-trained on the immigration subdataset, is fine-tuned with a limited number of unseen data to predict subframes. The size of abortion and gun data is limited to 50 and 100 triplets. For both topics, our model's accuracy and $F_1$ scores are higher than the baseline, when the data size is 50 and 100 triplets. This result shows that our pre-trained model only requires a small amount of labeled data for fine-tuning, which improves the scalability of data annotation.

In order to understand the influence of political bias on framing, we compute the average differences in subframe usage between biases. Our contrastive objective encourages CLoSE to separate different subframe groups farther apart. Simultaneously, binary political bias information is given as context to the embedding space. We would like to verify whether the bias objective distinguishes amongst subframes that are used by both ideologies in similar usage percentages. For each subframe group, we measure the percentage of its usage by the liberal and conservative media sources and compute the average of the difference in those

percentages. We then use the subframe group classifications of CLoSE and compare them to the true subframe group labels of the original data.

The results are shown in Table 8. The average differences of our model were *smaller* than those of the original data across the datasets. Namely, the overlapping usage of subframes between ideologies is more clearly observed in our model.

Figures 3a and 3b are the visualizations of the embeddings of the immigration data. Principal Component Analysis (PCA) IS applied to reduce dimensions, and the reduced embeddings ARE plotted and labeled with color. Figure 3a displays the clustering of subframe groups, and Figure 3b shows the distribution of the liberal and conservative biased embeddings. Importantly, Figure 3b shows the intersection of subframes between political bias, which aligns with the experimental results of Table 8. The embedding visualizations of the abortion and gun control data can be found in Appendix C.

### 6.4 Topic Modeling for Frame Extension

Lastly, we propose a topic modeling method that leverages the subframe embeddings output by our proposed CLoSE model to predict new indicators which can be used to extend or generalize the subframes. First, we cluster the framing embeddings, the outputs of our model, with $k$-means clustering. Then, we use class-based TF-IDF to cluster the framing embeddings and to find topics (Grootendorst, 2022). TF-IDF is a classic keyword extraction method that combines term frequency and inverse document frequency (Joachims, 1996). The TF-IDF score of a token $t$ in a document $D$ is:

$$W_{t,D} = tf_{t,D} \cdot \log \frac{N}{df_t},$$

where $tf_{t,D}$ is the frequency of the token $t$ in the document $D$, $df_t$ is the total number of documents that contain $t$, and $N$ is the total number of documents. The class-based TF-IDF generalizes TF-IDF to clusters by considering all documents in a cluster as a single document. It is defined as:

$$W_{t,C} = tf_{t,C} \cdot \log(1 + \frac{A}{tf_t}),$$

where $tf_{t,C}$ is the frequency of the token $t$ in a cluster $C$, $tf_t$ is the frequency of $t$ across all clusters, and $A$ is the average number of tokens per cluster.

Finally, we merge the clusters with high overlap of keywords, or tokens with high class-based TF-IDF scores. Suppose there are $c$ clusters $C = \{C_1, C_2, \cdots, C_c\}$. For a cluster $C_i$, we rank tokens according to their class-based TF-IDF scores and choose the top-$k$ number of tokens. Then, we compute the overlap of the tokens across clusters:

$$s(C_i, C_j) = \frac{\text{count}(C_i \cap C_j)}{|C_i|},$$

where $i \neq j$. If $s(C_i, C_j)$ is greater than a threshold $\beta$, we merge $C_i$ and $C_j$. For our experiments, the default values are $k = 50$ and $\beta = 0.4$.

| | |
|---|---|
| 1 | reproductive, health, justice, women, freedom |
| 2 | federal, funding, abortions, Hyde, Amendment |
| 3 | ⋆ |
| 4 | pregnancy, crisis, centers, Pro-Life, women |
| 5 | Wade, Roe, overturn, supreme, court |
| 6 | rights, anti-abortion, religious, catholic, abortion |
| 7 | Affordable, Care, Act, health, insurance |
| 8 | control, birth, health, save, prescription |
| 9 | Planned, Parenthood, selling, unborn, videos |
| 10 | life, unborn, right, baby, child |
| 11 | March, Life, Washington, national, Trump |
| 12 | industry, abortion, giant, business, profit |
| 13 | fetal, tissue, research, illegally, selling |
| 14 | Hobby, Lobby, access, coverage, insurance |

⋆ Potentially offensive or triggering language has been omitted.

Table 9: Topics of Subframe Embeddings for Abortion.

Table 9 shows the experimental results on the abortion data. We observe the subframe indicators of the original dataset as keywords. For instance, topic 9 has the subframe indicators: "Planned Parenthood." Interestingly, our topic modeling method is able to identify important words that are *not* in the subframe lexicons. Topic 2 has known indicators from the subframe "Abortion Funding" as well as new indicators "Hyde" and "Amendment."[3] This indicates CLoSE's ability to identify new subframe indicators which would allow the development of new frames or unsupervised frame detection. Appendix D presents experimental results for the topics and subdatasets of gun control and immigration.

## 7 Conclusion

We have presented CLoSE, a framework that incorporates subframes and political bias using contrastive learning and multi-task learning. Our proposed joint objective outperforms baseline models

[3]The Hyde Amendment, which was passed in 1976 by the House, is a legislative provision that prohibits the use of federal funds for performing abortions.

in predicting subframes and political bias. Further, our pre-trained model adapts to zero-shot learning of political bias and few-shot learning of subframes. Finally, we propose a topic modeling method for the subframe embeddings and extract a list of keywords, which can be helpful for future subframe extensions and annotations. We plan to extend this work to embed general policy frames for study in unsupervised settings.

## References


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.

Vibhu Bhatia, Vidya Prasad Akavoor, Sejin Paik, Lei Guo, Mona Jalal, Alyssa Smith, David Assefa Tofu, Edward Edberg Halim, Yimeng Sun, Margrit Betke, Prakash Ishwar, and Derry Tanti Wijaya. 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, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

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

Amber E Boydstun, Dallas Card, Justin Gross, Paul Resnick, and Noah A Smith. 2014. Tracking the development of media frames within and across policy issues.

Dallas Card, Amber E. 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, Beijing, China. Association for Computational Linguistics.

Kenneth Church and Patrick Hanks. 1990. Word association norms, mutual information, and lexicography. *Computational linguistics*, 16(1):22–29.

Carroll Doherty, Jocelyn Kiley, Nida Asheer, and Calvin Jordan. 2021. Beyond red vs. blue: The political typology. *Pew Research Center*.

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. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3570–3580, Brussels, Belgium. Association for Computational Linguistics.

Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.

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.

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. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4479–4488, Online. Association for Computational Linguistics.

Shanto Iyengar. 1994. *Is anyone responsible?: How television frames political issues*. University of Chicago Press.

Jakob D Jensen, Courtney L Scherr, Natasha Brown, Christina Jones, Katheryn Christy, and Ryan J Hurley. 2014. Public estimates of cancer frequency: cancer incidence perceptions mirror distorted media depictions. *Journal of health communication*, 19(5):609–624.

Thorsten Joachims. 1996. A probabilistic analysis of the rocchio algorithm with tfidf for text categorization. Technical report, Carnegie-mellon univ pittsburgh pa dept of computer science.

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, Osaka, Japan. The COLING 2016 Organizing Committee.

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, Vancouver, Canada. Association for Computational Linguistics.

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, Vancouver, Canada. Association for Computational Linguistics.

Grzegorz Kapuściński and Barry Richards. 2016. News framing effects on destination risk perception. *Tourism Management*, 57:234–244.

Johannes Kiesel, Maria Mestre, Rishabh Shukla, Emmanuel Vincent, Payam Adineh, David Corney, Benno Stein, and Martin Potthast. 2019. SemEval-2019 task 4: Hyperpartisan news detection. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 829–839, Minneapolis, Minnesota, USA. Association for Computational Linguistics.

Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

Bingjie Liu and Lori Pennington-Gray. 2015. Bed bugs bite the hospitality industry? a framing analysis of bed bug news coverage. *Tourism Management*, 48:33–42.

Amy Mitchell, Jeffrey Gottfried, Jocelyn Kiley, and Katerina Eva Matsa. 2014. Political polarization & media habits.

Vinod Nair and Geoffrey E Hinton. 2010. Rectified linear units improve restricted boltzmann machines. In *Icml*.

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.

Zhongdang Pan and Gerald M Kosicki. 1993. Framing analysis: An approach to news discourse. *Political communication*, 10(1):55–75.

Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.

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, Online. Association for Computational Linguistics.

Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.

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.

Isidora Tourni, Lei Guo, Taufiq Husada Daryanto, Fabian Zhafransyah, Edward Edberg Halim, Mona Jalal, Boqi Chen, Sha Lai, Hengchang Hu, Margrit Betke, Prakash Ishwar, and Derry Tanti Wijaya. 2021. Detecting frames in news headlines and lead images in U.S. gun violence coverage. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4037–4050, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Oren Tsur, Dan Calacci, and David Lazer. 2015. A frame of mind: Using statistical models for detection of framing and agenda setting campaigns. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1629–1638, Beijing, China. Association for Computational Linguistics.

Arjen Van Dalen, Claes de Vreese, and Erik Albæk. 2017. Economic news through the magnifying glass: How the media cover economic boom and bust. *Journalism Studies*, 18(7):890–909.

## Appendix A Subframes

Table 10 presents the list of subframes for each topic. As subframes are issue-specific subclassifications of the general policy frames, the frames are written in bold, and subframes that fall under a corresponding frame are displayed below the frame. This table is from Roy and Goldwasser (2020).

## Appendix B Top-3 Subframes

Table 11 shows the top-3 subframes used across the period for the topic of gun control. Similarly, Table 12 is the results for the topic of immigration. High overlapping of subframes used by liberal and conservative media is observed in both data. As for the gun control data, all top-3 subframes from the liberal and conservative articles from 2020 to 2022 are identical. Similarly, the top-2 subframes from the liberal and conservative news from 2020 to 2022 are the same for the immigration data. We did not extract the top-3 subframes of the immigration data before 2010 because all the articles in the immigration dataset were from liberal media sources.

## Appendix C Embedding Visualization

Figures 4 and 5 are the embedding visualizations of the abortion data and gun control data, respectively. The dimensions of the output embeddings of our proposed method are reduced with PCA, and the reduced embeddings are mapped to 2D plots. Subfigure (a) is labeled with color to show the clustering of subframes. Subfigure (b) displays the liberal embeddings in blue and the conservative embeddings in red. As we observed in the embedding plots of the abortion data, Figures 4b and 5b display a high percentage of overlap between the subframes used by liberal and conservative articles.

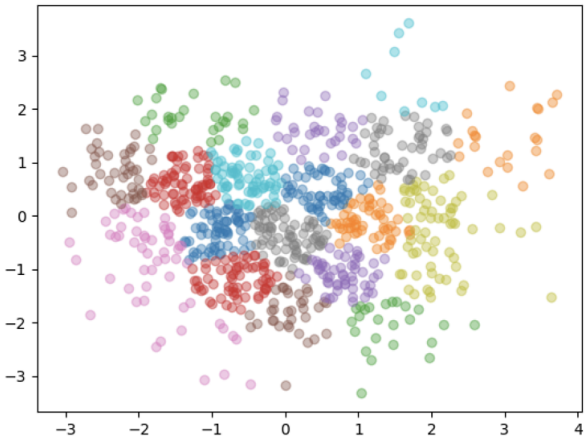


(a) Subframes

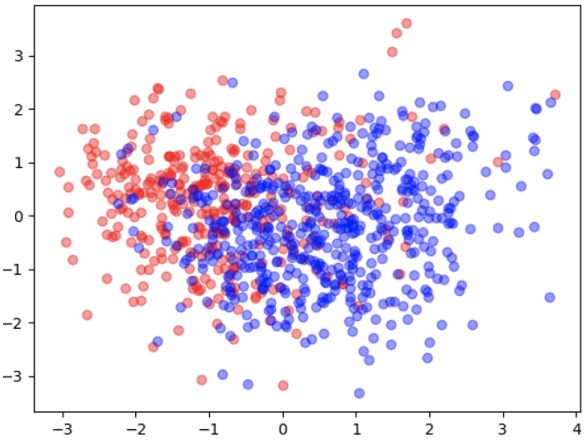


(b) Political Bias

Figure 4: Visualization of the Embeddings of the Abortion Dataset. Subfigure (a) shows the clustering of subframe groups. Subfigure (b) is labeled based on the political bias of each article.

## Appendix D Topic Modeling of Framing Embeddings

Table 13 shows the topics of the subframe embeddings of the gun control data, and Table 14 shows those for the immigration data. After merging the topics with high overlapping keywords, there are 12 and 15 topics for the gun control and immigration data, respectively.

As for the topics of the gun control data, names and locations related to mass shootings are observed in topics 10 and 11. In topic 6, the words "Islamic" and "York" from New York appear along

| Abortion | Gun Control | Immigration |
|---|---|---|
| **Economic**: | **Economic**: | **Economic**: |
| - Health Care | - Gun Buyback | - Minimum Wage |
| - Abort. Provider | Program | - Salary Stagnation |
| - Abortion Funding | - Gun Business | - Wealth Gap |
| **Fairness & Equality:** | **Capacity & Resources:** | - Cheap Labor Availability |
| - Reproduction Right | - School Safety | - Taxpayer Money |
| - Right of Human Life | **Cultural Identity:** | **Crime & Punishment:** |
| **Legality, Constitution** | - White Identity | - Deportation: Illegal |
| **ality, Jurisdiction:** | - Person of Color Identity | Immigrants |
| - Hobby Lobby | **Legality, Constitution** | - Deportation: In General |
| - Late Term Abortion | **ality, Jurisdiction:** | - Detention |
| - Roe V. Wade | - Ban on Handgun | - Terrorism |
| **Crime & Punishment:** | - Second Amendment | **Security & Defense:** |
| - Stem Cell Research | - Concealed Carry | - Border Protection |
| - Sale of Fetal Tissue | Reciprocity Act | **Legality, Const., Juri.** |
| - Sexual Assault Victims | - Gun Control to | - Asylum |
| **Health & Safety:** | Restrain Violence | - Refugee |
| - Birth Control | **Crime & Punishment:** | - Birth citizenship & |
| **Morality:** | - Illegal Gun | 14th Amendment |
| - Sanctity of Life | - Gun Show Loophole | **Policy Pres. & Eval.:** |
| - Women Freedom | **Security & Defense** | - Amnesty |
| **Quality of Life:** | - Background Check | - Dream Act |
| - Planned Parenthood | - Terrorist Attack | - Family Separation |
| - Pregnancy Centers | **Health & Safety:** | Policy |
| - Life protection | - Gun Research | - DACA |
| **Public Sentiment:** | - Mental Health | **Fairness & Equality:** |
| - Pro-Life | - Gun Homicide | - Racism & Xenophobia |
| - Anti-Abortion | **Policy Pres. & Eval.:** | - Merit Based Immigration |
| - Pro-Choice | - Assault Weapon | - Human Right |
| | **Morality:** | **Cultural Identity:** |
| | - Right to Self-Defense | - Racial Identity |
| | - Stop Gun Crime | - Born identity |

Table 10: Subframes of the Three Topics: Abortion, Gun Control and Immigration. Frames are written in bold.

| | | Top-1 | Top-2 | Top-3 |
|---|---|---|---|---|
| 2020-2022 | Left | Gun Business Industry | Conc. Carry Recip. Act | Second Amendment |
| | Right | Gun Business Industry | Conc. Carry Recip. Act | Second Amendment |
| 2018-2020 | Left | Background Check | School Safety | Assault Weapon |
| | Right | Background Check | Assault Weapon | Second Amendment |
| 2016-2018 | Left | Gun Business Industry | Terrorist Attack | Conc. Carry Recip. Act |
| | Right | Gun Business Industry | Terrorist Attack | Assault Weapon |
| 2010-2016 | Left | Gun Business Industry | Background Check | Mental Health |
| | Right | Terrorist Attack | Gun Show Loophole | Gun Business Industry |
| -2010 | Left | Second Amendment | Gun Business Industry | Terrorist Attack |
| | Right | Gun Business Industry | Illegal Gun | Second Amendment |

Table 11: Top-3 Subframe Indicators of the Gun Control Triplet Dataset.

| | | Top-1 | Top-2 | Top-3 |
|---|---|---|---|---|
| 2020-2022 | Left | Asylum | Detention | Human Rights |
| | Right | Asylum | Detention | Racial Identity |
| 2018-2020 | Left | Asylum | DACA | Human Rights |
| | Right | Asylum | DACA | Amnesty |
| 2016-2018 | Left | DACA | Human Rights | Racial Identity |
| | Right | DACA | Amnesty | Dream Act |
| 2010-2016 | Left | DACA | Human Rights | Detention |
| | Right | Birth Cit. & $14^{th}$ Amen. | DACA | Dream Act |

Table 12: Top-3 Subframe Indicators of the Immigration Triplet Dataset.

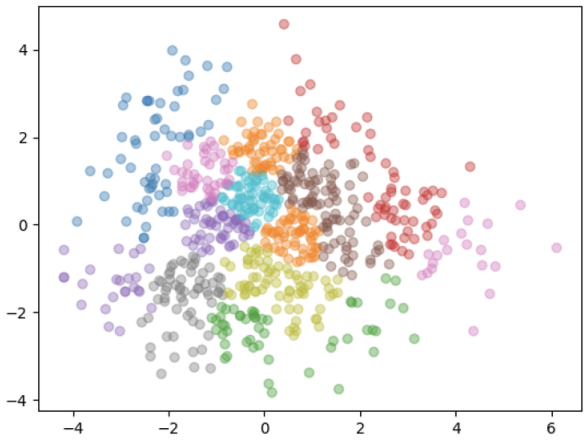

(a) Subframes

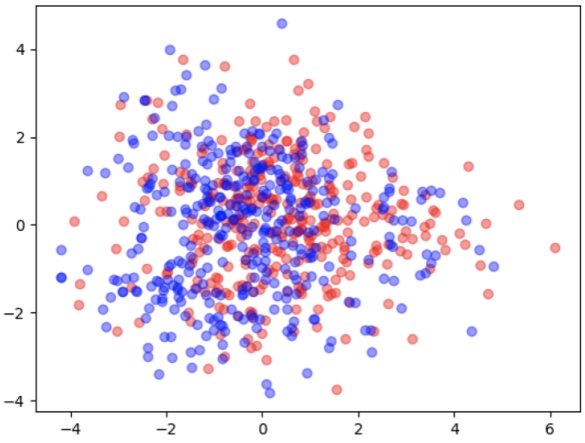

(b) Political Bias

Figure 5: Visualization of the Embeddings of the Gun Control Dataset. Subfigure (a) shows the clustering of subframe groups. Subfigure (b) is labeled based on the political bias of each article.

with the subframe indicators related to "terrorism," which imply the co-ocurrence of a specific religion with a place.

Similarly, Table 13 shows specific names and places such as "Germany," "Mexico," and "ACLU," which is an acronym for the American Civil Liberties Union that fights for civil rights, including the rights of immigrants. The names of the two former presidents are also found: "Trump" in topic 14 along with the keywords "build" and "wall" and "Obama" in a keyword "Obama-era" in topic 5 with keywords associated with the Dream Act.

| | |
|---|---|
| 1 | Amendment, Second, rights, protect, individual |
| 2 | background, checks, NICS, FBI, criminal |
| 3 | industry, manufacturers, business, selling |
| 4 | rifles, gun, firearms, semiautomatic, automatic |
| 5 | illegal, guns, possessions, convicted, dealers |
| 6 | terrorism, terrorist, Islamic, York, Trump |
| 7 | ill, mentally, illness, violence, mental |
| 8 | abusers, shootings, victims, health, age |
| 9 | ban, handgun, Concealed, Carry, Reciprocity |
| 10 | school, violence, students, Parkland, Florida |
| 11 | contempt, Orlando, Mateen, Black, Americans |
| 12 | Buyback, program, deaths, restrictions, mandatory |

Table 13: Topics of Subframe Embeddings of the Gun Control Data.

| | |
|---|---|
| 1 | detention, center, facilities, ICE, children |
| 2 | Amnesty, plan, seekers, refugee, Germany |
| 3 | Families, Separation, Policy, Trump, parents |
| 4 | DACA, Democrats, childhood, arrivals, program |
| 5 | Act, Dream, dreamers, children, Obama-era |
| 6 | asylum, migrants, foreign, illegal, caravan |
| 7 | rights, human, Civil, Liberties, Union, ACLU |
| 8 | black, brown, undocumented, immigrants, racial |
| 9 | labor, cheap, workers, wages, manual |
| 10 | illegal, deported, alien, deportation, persons |
| 11 | white, identity, supremacist, nationalist, racist |
| 12 | terrorist, threat, border, terrorism, terror |
| 13 | birthright, citizenship, Amendment, 14th, Anton |
| 14 | build, wall, Trump, Republican, Mexico |
| 15 | minimum, wage, inequlaity, income, raise |

Table 14: Topics of Subframe Embeddings of the Immigration Data.

## Appendix E Ethical Considerations

The Framing Triplet Dataset is an extension of the existing public dataset by Kiesel et al. (2019) and Roy and Goldwasser (2020). Additionally collected documents are also from news texts that are free to the public. Hence, the text corpus does not contain private or sensitive information.

The code for the proposed method is open to the public and can be used to study the impact of news framing on public opinion. We do not anticipate any significant risks of deployment. Still we urge users not to use this research for malicious intentions.