



POLOR: Leveraging Contrastive Learning to Detect Political Orientation of Opinion in News Media

Ala Jararweh and Abdullah Mueen

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

May 12, 2024

POLOR: Leveraging Contrastive Learning to Detect Political Orientation of Opinion in News Media

Ala Jararweh

The University of New Mexico
Albuquerque NM, USA
ajararweh@unm.edu

Abdullah Mueen


The University of New Mexico
Albuquerque NM, USA
mueen@unm.edu

Abstract

News articles are naturally influenced by the values, beliefs, and biases of the reporters preparing the stories and the policies of the publishing outlets. Numerous studies and datasets have been proposed to detect the political orientation of news articles. However, most of these studies ignore real textual clues and learn the textual signature of the source (commonly the publisher and rarely the writer) of the article instead. Moreover, a good volume of opinion pieces published by major news outlets do not reflect the political orientation of the publisher but rather reflect the political orientation of a non-professional writer. Existing methods are not built to correct this difference in the training data and, hence, perform poorly on human-annotated data. We propose, POLOR, a fine-tuned BERT model that employs contrastive learning to detect the political orientation of news articles even when the training data is labeled by the source (i.e. the publisher of the news article). Unlike previous work in the literature, the model learns features by employing different contrastive learning objectives where each sentence is contrasted with sentences from various sources simultaneously. POLOR achieves a 15% increase on our dataset compared to previously proposed baselines. Finally, we release two datasets of opinion news: source-annotated and human-annotated datasets. The full paper including supplementary materials, code, and datasets can be found at <https://www.cs.unm.edu/~ajjararweh/>.

1 Introduction

Political bias in news can be emphasized in various forms such as rephrasing, or presenting one-sided facts to serve a specific orientation. This bias directly impacts the information consumed by the readers and the political attitudes (Gentzkow and Shapiro 2006). Detecting the political orientation of news is a challenging task due to the necessity of understanding the multidimensional political discourse. For example, mainstream media worldwide use terms like "explosion" versus "attack" when reporting on recent conflicts in the Middle East to favor one political position over another (de Jong 2023). This issue has



(a) Crowd-Sourcing Labels

(b) Source Labels

Figure 1: Visualized article embeddings generated by POLOR, and annotated according to (a) article orientation, and (b) source orientation. The blue points represent *Liberal*, and the red points represent *Conservative*. Source annotations do not always reflect all article orientations published by that source. Consequently, models trained on source annotations without utilizing auxiliary information, learn to predict sources instead of capturing article-specific orientations in Figure 1(a).

been studied in various contexts and a variety of datasets has been proposed. Researchers typically use online platforms such as *AllSides.com*, *AdFontesMedia.com*, *MediaBiasFactCheck.com*, etc to collect and annotate news articles.

Such platforms usually offer political annotations at the news source level. To obtain article-specific orientation, annotations are usually propagated from the source annotations (Lee et al. 2022; Baly et al. 2020; Chen et al. 2018; Liu et al. 2022; Kulkarni et al. 2018). For example, *AllSides.com* annotates all articles published by the New York Times as *Liberal*. Figure 1 shows visualized embeddings of articles generated by our model. The dimensions of the embeddings were reduced using t-distributed Stochastic Neighbor Embedding (van der Maaten and Hinton 2008). The true article orientations (Figure 1(a)) differ from the annotations derived from sources (Figure 1(b)). The red points are more distributed to the right and the blue points are to the left in both figures which illustrates that the source orientations can be related to the article orientations.

Relying on the source annotation without utilizing auxil-

ary knowledge, the trained models may attempt to learn the writing style of the news source instead of predicting the orientations of news articles. The problem is emphasized when opinion pieces are considered. Opinion pieces usually represent the bias of the authors, and are often published by a media with opposite bias. This problem has been lightly touched on by Baly et al. (2020), where the authors found that the model suffers when trained on source-labeled data and tested on unseen data (i.e. new media sources). To overcome this, they inserted a source classifier to minimize the knowledge learned from the source. However, they have not tested their model on human-annotated data where the labels are derived based on the text in articles, paragraphs, or sentences.

In this paper, we propose **POLOR**: leveraging contrastive learning to detect the **POL**itical **OR**ientation of opinion pieces in news media. POLOR exploits several contrastive learning objectives where each sentence is contrasted to a set of sentences from different and similar sources. We adopt different objective functions to generate additional features that focus on textual cues related to the political bias of articles, instead of the style of the news media. POLOR produces sentence-specific and article-specific labels based on the training data derived from sources.

2 Related Work

Detecting the political orientation in news media can be challenging due to the necessity of domain expertise. In their early attempts, researchers measured the orientation of a particular news media by counting the number of citations to liberal and conservative think-tank policy groups (Groseclose and Milyo 2005). Focusing on the article content, Greene and Resnik (2009) used domain-specific words in news articles to classify articles that support Palestine or Israel. Similarly, researchers harnessed the recent findings in Deep Learning to understand the political orientation of news articles at different levels of article structure such as character-based word representations (Jiang et al. 2019), sentences (Kim and Johnson 2022; Gangula, Duggenpudi, and Mamidi 2019; Chen et al. 2018), and articles (Baly et al. 2020; Chen et al. 2020; Liu et al. 2022).

Various datasets were proposed to tackle this task with a variety of annotation techniques. Since it’s challenging and expensive to annotate a large corpus of news articles, different studies rely on online platforms such as `AllSides.com` to crawl and annotate news articles (Baly et al. 2020; Chen et al. 2020; Liu et al. 2022; Roy and Goldwasser 2020). The annotations of articles are usually propagated from the source orientation. However, a major drawback of these annotations is that models learn to predict the source signature rather than the article’s orientation. To overcome this, researchers use online crowd-sourcing platforms to obtain article-specific annotations (Gangula, Duggenpudi, and Mamidi 2019; Fan et al. 2019; Card et al. 2015; Budak, Goel, and Rao 2016; Lazaridou et al. 2020). To address the problem of predicting article-specific labels from source-based annotations, few datasets were found that integrate both annotations. Kiesel et al. (2019) proposed two datasets but the human annotation dimensions of articles are

different from our task. To this end, we construct a new dataset that spans 10 news sources with two different annotations.

	Liberal	Conservative	
Articles#	156	107	263
paragraphs#	5146	3099	8245
Sentences#	7743	3920	11921
Sources#	4	3	7

(a) Statistics of the Source-Based dataset.

	Liberal	Conservative	Neutral	
Articles#	22	17	0	39
Paragraphs#	138	133	42	313
Sentences#	522	513	188	1223
Sources#	-	-	-	3

(b) Statistics of the Human-annotated dataset.

Table 1: Exploring the constructed dataset for this task.

3 Datasets

We evaluate our model performance on two datasets. The first dataset consists of approximately 6,000 articles annotated via crowd-sourcing from Budak, Goel, and Rao (2016). It also spans 11 news sources and explores 15 different topics. Since the source annotations are not available, we extend the dataset by deriving annotations from `AllSides.com`. For a broader evaluation, we construct another dataset for this task. Our dataset consists of two subsets, crowd-sourcing annotations and source annotations. The constructed dataset addresses two high-profile criminal cases, namely *Ahmaud Arbery*¹ and *Kyle Rittenhouse*².

Source-Based Dataset. We collected a total of 263 news articles which cover 7 different U.S. news sources. We align with recent work in the literature, we obtained annotations from `AllSides.com`. This website only offers annotations at the source level based on multiple methods such as Editorial Review, Blind Bias Survey, and Community Feedback. We collect articles from 4 liberal news sources and 3 conservative news sources according to the annotations found on the website. Table 1(a) shows some detailed statistics of the dataset. We follow the same annotation process of recently published datasets where the article annotations are propagated from their source annotations.

Human-Annotated Dataset. We also construct a smaller dataset for evaluating the model on human annotations. This dataset consists of 39 news articles from three alternative U.S. news sources, different from those in the training data (Source-Based Dataset). We aim to evaluate our model’s ability to generalize to new writing styles. Thus, we only consider opinion news published about the two cases. The articles were collected via word matching from the news

¹https://en.wikipedia.org/wiki/Murder_of_Ahmaud_Arbery

²https://en.wikipedia.org/wiki/Kyle_Rittenhouse

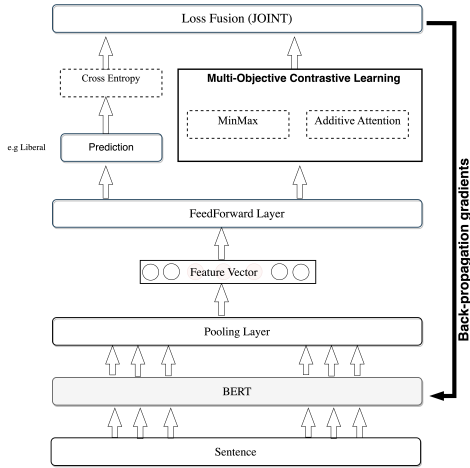


Figure 2: An overview of POLOR, which utilizes contrastive learning objectives to extrapolate article-specific labels.

media websites. The annotations were performed at the paragraph level, where each article was partitioned into a set of paragraphs to provide annotators with more context about the cases. The article annotations are then derived based on majority voting. Each paragraph is labeled by three different workers into one of three different categories: Liberal, Conservative, or Neutral. The workers were provided with a set of instructions and facts about the two cases ahead of time. 92% of the articles received an inter-annotator score above 0.9, and 8.0% are below 0.9. The articles that have low scores (below 0.5) were annotated Neutral. The news source of each paragraph was hidden from the workers and they only could see the content of the paragraph. Table 1(b) shows the detailed statistics and the annotation results. To match with annotation dimensions in the source data, we only consider conservative and liberal paragraphs in our evaluation.

Data Preparation. The datasets are further pre-processed by removing unnecessary content such as punctuation, non-English characters, and identifying words and sentences. We then split each article/paragraph into a set of sentences. The final preparation step is to formulate the datasets as triplets. Each sentence in the dataset S_a (also called Anchor) is attached to a set of positive sentences S_p , and a set of negative sentences S_n . We define a sentence as positive if it belongs to the same class and from a different source, and negative if it belongs to the opposite class. For example, an anchor labeled as "Conservative", its positive set is randomly derived from other sentences labeled as "Conservative", and its negative set is randomly derived from sentences labeled as "Liberal". We also ensured that the triplets of the training, validation, and test set were completely disjoint. The size of the positive and negative sets (\mathcal{T}) is a parameter.

4 Methodology

We propose POLOR, a complete framework that leverages multi-objective contrastive learning to generate embeddings

that portray the bias in news articles by minimizing the source information. Figure 2 shows the main components of the model, namely Sentence Embedding, Multi-Objective Contrastive Learning, Prediction, and Loss Fusion (Joint Loss).

4.1 Sentence Embedding

The sentences are first tokenized, where each sentence in anchor, negative, or positive sets is represented as a list of N tokens. The tokens are passed to the pre-trained BERT model to generate contextualized token representations (i.e. embeddings) where each token, T_i , is represented in an embedding vector, E_i . Figure 2 shows the input and output of BERT. At the end of this stage, each sentence S will be represented with N contextualized token representations, where $E_i \in \mathbb{R}^{768}$. To synthesize sentence embeddings, we use the pooling strategies, CLS (classification token) Pooling and Mean Pooling introduced by Reimers and Gurevych (2019).

4.2 Multi-Objective Contrastive Learning

Our goal is to tune the representations such that they exclusively represent the bias in the content of articles. Since the training dataset (Source-Based Dataset) only contains sentences annotated based on sources, the model learns additional features in an unsupervised manner. That is, the source influence is minimized and the influence of the political orientation in the text is maximized. To achieve this goal, we employ different objective functions that utilize contrastive learning. In the results section, we empirically show that augmenting contrastive objectives on top of BERT helps to learn rich features by contrasting other sentences from similar/different sources in the dataset. We attempt to find a better representation by employing a triplet loss function. The triplet loss function was introduced by Schroff, Kalenichenko, and Philbin (2015) in the field of face recognition to optimize face image embeddings. The authors define positive examples as similar images of the same face, and negative examples are all other faces in the mini-batch.

We adopt an alternative solution to this. We first assign a set of \mathcal{T} positives and a set of \mathcal{T} negatives from the entire dataset at random for each anchor. We then apply different objective functions on the positive and negative sets to synthesize rich representations to achieve different optimizations. After applying each objective function, the synthesized representation should have a similar dimension to the anchor (\mathbb{R}^{768}). The synthesized representations of positive and negative sets are then contrasted with the anchor to obtain the loss for the desired objective. Given an anchor sentence S_a , set of positive sentences S_p and, and set of negative sentences S_n , where S_p and S_n of size \mathcal{T} , we re-frame the loss function as follows:

$$\mathcal{L}(S_a, S_p, S_n) = \max \begin{cases} \|S_a - \mathcal{F}(S_p; 1, \mathcal{T})\|_2 - \|S_a - \mathcal{F}(S_n; 1, \mathcal{T})\|_2 + m \\ 0 \end{cases} \quad (1)$$

where m is the margin to be enforced between the positive and negative sets, and $\mathcal{F}(S; 1, \mathcal{T})$ is an objective function used to synthesize a representation from the set of positive and negative sets. In the following subsections, we describe the objective functions used to synthesize the representations. We then define the Joint loss function which is designed to fuse all losses, obtained by contrasting data to the synthesized representations, into one universal function.

Additive Attention. Since the sentences are annotated based on source, the annotation does not always reflect the true orientation of the sentences. The model attends to the sentences in the positive and negative sets based on importance and relativity to the anchor. We modify the additive attention mechanism mentioned in Bahdanau, Cho, and Bengio (2014) to serve our goal. This attention mechanism was mainly proposed in machine translation to synthesize a context vector of words in a sentence. The context vector is a weighted sum of the encoder’s hidden state with the attention scores. The attention scores are usually computed by feeding the concatenation of the encoder’s current hidden state and the decoder’s previous hidden state to an alignment model (a non-linear transformation of the input).

In our case, the position of sentences in the positive/negative set does not matter in predicting the anchor. Consequently, we replace the hidden state of the encoder and the decoder with two feed-forward layers where the input is the positive/negative sentence embeddings. The results are then concatenated and fed to an alignment model to calculate the attention scores.

$$e = W_3 \cdot \tanh([W_1 \cdot E; W_2 \cdot E])$$

$$a_i = \frac{\exp(e_i)}{\sum_{k=1}^{\tau} \exp(e_k)} \quad (2)$$

where E is the embedding of sentences from the positive/negative set. The final step is to calculate the weighted sum of positive/negative embeddings with the attention scores:

$$\mathcal{F}(E; 1, \mathcal{T}) = \sum_{i=1}^{\tau} a_i \cdot E_i \quad (3)$$

where we obtain $\mathcal{L}_{\mathcal{F}_{additive}}$ by substituting $\mathcal{F}(E; 1, \mathcal{T})$ in Eq. 1 for the negative and positive sets separately.

Unsupervised MinMax. The goal of this objective function is to find the appropriate negative and positive sentences extracted from the entire positive and negative sets. We define a sentence as appropriate such that it influences the model to increase the margin between the negative set and the positive set in the Euclidean space. To achieve this goal, unsupervised MinMax treats the positive and negative sentences as one set where the closest sentence is considered negative and the furthest sentence is considered positive.

$$\mathcal{F}(S_p; 1, \mathcal{T}) = \operatorname{argmax}_{s \in (S_p \cup S_n)} \|g(S_a) - g(s)\|_2^2$$

$$\mathcal{F}(S_n; 1, \mathcal{T}) = \operatorname{argmin}_{s \in (S_p \cup S_n)} \|g(S_a) - g(s)\|_2^2 \quad (4)$$

where $g(s)$ is the embedding of the sentence s . The loss function $\mathcal{L}_{\mathcal{F}_{MinMax}}$ is then obtained by substituting in Eq. 1.

This selection ensures fast convergence and helps to increase the margin between the anchor and negative sentences and decrease it with positive sentences (Schroff, Kalenichenko, and Philbin 2015).

4.3 Joint Loss (Fusion)

The losses obtained by the above objective functions are assembled to allow the model to learn better representation. We use different hyper-parameters to weigh the objective losses and then join them into one loss function. In addition to the unsupervised contrastive loss, we incorporated Binary Cross-Entropy loss (\mathcal{L}_{XEnt}) in the final loss function to obtain the Joint loss (\mathcal{L}_J):

$$\mathcal{L}_J = \alpha \cdot \mathcal{L}_{\mathcal{F}_{additive}} + \beta \cdot \mathcal{L}_{\mathcal{F}_{MinMax}} + (1 - \alpha - \beta) \cdot \mathcal{L}_{XEnt} \quad (5)$$

where $\mathcal{L}_{\mathcal{F}}$ is the triplet loss obtained by applying a specific objective function, and $0 < \alpha, \beta < 0.5$. Note that, when a parameter equals 0, the associated loss function is not incorporated into the Joint loss.

5 Experiments and Results

5.1 Baselines

Our goal is to classify the article-based political orientation of an article using training data labeled by their sources. We evaluate our model performance against different baselines with different training setups of our model:

1. **Majority classifier:** We consider the most frequent class in the dataset as the predicted label for all sentences.
2. **Joint function with XEnt loss only :** We use the sentence representations generated by BERT to perform classification in a supervised manner. For this baseline, the contrastive learning loss is entirely ignored (i.e. $\alpha = 0$ and $\beta = 0$). This approach is similar to Fan et al. (2019).
3. **Joint function with XEnt and Additive:** We considered two loss functions in the joint function (Eq. 5), namely additive and XEnt. That is, β is set 0.
4. **Joint function with XEnt and MinMax:** We consider two loss functions in the joint function (Eq. 5), namely XEnt and MinMax. That is, β is set 0.
5. **Joint function with XEnt loss and Random Triplet objective:** We consider only one contrastive objective function that selects a random positive and negative sentence from the sets S_p and S_n at each iteration. The joint loss becomes:

$$\mathcal{L}_J = \alpha \cdot \mathcal{L}_{\mathcal{F}_{RT}} + (1 - \alpha) \cdot \mathcal{L}_{XEnt}$$

This setup is approximately similar to Kim and Johnson (2022) and Baly et al. (2020) with few differences. In their formulation, the triplet is usually chosen at random in an offline manner (i.e. when the dataset is created) and remains fixed throughout the whole training process. Moreover, their method of choosing positive and negative sentences is different from our approach. For example, Kim and Johnson use predefined sub-framing groups to assign the triplets, while Baly et al. use media sources.

Losses	Model	Large Media		Our dataset			
		Acc _{Art}	F1 _{Art}	Acc _{Sent}	F1 _{Sent}	Acc _{Art}	F1 _{Art}
Majority	-	0.52	0.34	0.50	0.34	0.51	0.34
XEnt (BERT)	Fan et al. (2019)	0.63	0.60	0.51	0.51	0.62	0.61
XEnt + Additive	ours	0.65	0.62	0.53	0.53	0.67	0.65
XEnt + MinMax	ours	0.63	0.60	0.53	0.53	0.69	0.68
XEnt + Random	Kim and Johnson (2022), Baly et al. (2020)	0.62	0.59	0.53	0.47	0.64	0.64
XEnt + Centroid	ours	0.63	0.61	0.53	0.53	0.64	0.62
XEnt + Additive + MinMax	POLOR	0.68	0.66	0.55	0.520	0.74	0.74

Table 2: The results of comparing our multi-objective model (POLOR) with the baselines. The results on our dataset are reported on sentence predictions (Sent) and article predictions (Art) using accuracy (Acc) and macro F1 score metrics. For the large media study, the model was directly trained on articles.

6. **Joint function with XEnt loss and Centroid Triplet:** We consider only one contrastive objective function which is the mini-batch mean. In this objective, the positive is the centroid of the positive set and the negative is the centroid of the negative set. The joint loss becomes:

$$\mathcal{L}_J = \alpha \cdot \mathcal{L}_{\mathcal{F}_{centroid}} + (1 - \alpha) \cdot \mathcal{L}_{XEnt}$$

5.2 Quantitative Results

Our Dataset. Table 2 shows the performance of our model in predicting political orientation compared to the proposed baselines. POLOR outperforms all baseline models on sentence accuracy, article accuracy, and article macro F1 score. We use majority voting (highest occurring class) on sentences to predict the article labels. We first notice that the baseline models outperform the majority classifier when predicting the orientation of articles, but struggled with sentence orientation. Because sentence labels are propagated from articles, which are subsequently propagated from the sources. Ideally, sentence-specific labels would solve the problem. However, in the absence of sentence-specific labels, article-specific labels from independent sources would improve the performance. Moreover, choosing the proper objective function can drastically affect performance. Sophisticated objective functions may not always yield the desired performance. For example, the simple random triplet objective shows a competitive performance even though it requires almost no additional computations beyond selecting a sentence at random. Finally, additive and MinMax achieve a good performance, and when combined (POLOR), they achieve the highest outcomes. The combined objectives also illustrate a good performance on unseen source signatures, since the news sources in the testing set were not seen during training.

Large Media Dataset. The main objective of this experiment is to analyze the model behavior on a larger dataset. We evaluate our model performance on a dataset of 6089 news articles from Budak, Goel, and Rao (2016). The dataset is only annotated via crowd-sourcing. To obtain source annotations, we manually annotate articles based on source annotations from ALLSIDE.COM website. We test our model

performance on the entire article text to reduce the runtime. Table 2 shows the experiment results on this dataset. Our model of the combined Additive and MinMax objectives is still outperforming the baselines by a good margin. We also note that the additive objective separately achieves a competitive performance to POLOR.




Figure 3: Critical difference diagram showing pairwise statistical difference comparison of the proposed baselines.

Significance Test. To further explain POLOR performance compared to the proposed baselines, we conduct a critical difference test based on the Wilcoxon-Holm method to detect pairwise significance. The test was performed using 15-fold cross-validation on the large media dataset only. We have not performed the test on our dataset since the annotated test set is fixed, and running the experiment k-times would always yield the same results. Figure 3 shows the results of the significance test with a significance level of 0.05. POLOR ranks first compared to the proposed baselines which demonstrated a significant performance in predicting human-level labels from news source annotations. The test also illustrates the superior capability of the additive objective in this task, unlike the MinMax objective which is anticipated to be impacted by the distances between the articles in the positive and negative sets.

5.3 Qualitative Results

Visual Analysis of Model Attention. We further demonstrate the model performance by delving deeper into two correctly predicted sample paragraphs where the source annotations contradict the text annotations. Figure 4 depicts the model attention using the SHAP explainer (Lundberg and Lee 2017) where the conservative and liberal contributions are reflected by the red and blue highlights respectively. Although the selected examples are intricate and potentially

Barack Obama was thrown by the election of Scott Brown, a Republican, to fill a vacant Senate seat in heavily Democratic Massachusetts. The election not only cost Democrats their filibuster-proof supermajority just as they were trying to pass health care legislation, but it also foreshadowed a Republican landslide in midterm elections later that year.

(a) New York Times (Liberal), Paragraph Label: Conservative

We've endured too many of these tragedies in the past few years, Mr. Obama said, wiping away tears: The comments he made Friday on the need to act against gun violence mirror those he made in the aftermath of the 2011 shooting in Tucson, that severely wounded Rep. Gabrielle Giffords and killed six.

(b) Wall Street Journal (Conservative), Paragraph Label: Liberal

Figure 4: Visualizing the model attention on two real examples where the text label contradicts the source label.

elusive for some readers to catch the subtle undertone, the model efficiently allocated more attention to underlying signaling words and phrases. In Figure 4(a), the model demonstrates higher attention toward the term "Republican" compared to the term "Democrats". Similarly, the phrase "their ... care legislation" also contributes to the paragraph prediction since it is proposed or supported by a democratic group as the paragraph states, and the phrase "foreshadowed a Republican" which indicates a positive framing of anticipated Republican success. For Figure 4(b), the model attended to framing words that tend to humanize former President Barack Obama such as "wiping away tears" where it falls short in capturing other phrases like "act against gun violence". Since the paragraph labels exclusively contradict the source labels, the given remarks qualitatively suggest that the model devotes more attention to the signaling phrases in the text rather than learning the source writing style.




Figure 5: The performance of POLOR on unseen topics. The model shows a relatively stable performance and always outperforms the majority classifier even when we increase the number of hidden topics.

Cross-Topic Analysis. The baseline results in the quantitative section demonstrate that the model learns the political ideology of the news, not the source writing style since the test set was collected from unseen sources. However, that experiment falls short of explaining whether the model learns other textual properties such as topics. In this experiment, we are interested in measuring the model's generalizability on novel topics where we evaluate the model performance

on unseen topics during testing. We iteratively increase the number of unseen topics (from 0 to 5 topics). We report the accuracy values across five different runs to account for the variations and uncontrollable randomness. For comparability reasons, we also include the majority classifier accuracy on the given split. We perform this experiment on the Large Media dataset since it consists of multiple topics. Figure 5 shows the results of this experiment. The model shows relatively stable performance even when we increase the number of hidden topics. The model always outperforms the majority classifier even when we increase the number of hidden topics. That is, the model is still capable of predicting article-specific labels on novel topics that were not seen during training.

	Majority	Acc	F1
Source annotations	0.74	0.91	0.91
Crowd-sourcing	0.51	0.68	0.66

Table 3: Comparing article-specific orientation with source-specific orientation, using the same test data but annotated in two different ways.

Benchmarking Task Difficulty. Extrapolate article-specific orientation from source annotations is less straightforward than predicting source-specific orientation. In this section, we benchmark the two tasks. For both tasks, the training data is propagated from source annotations. However, for the article-specific task, the test data is annotated via crowd-sourcing, while the test data for the source-specific task continues to be annotated via source annotations. Table 3 shows how predicting source-specific orientations of articles yields good performance because the model learns to map the predictions to sources rather than predicting the article orientation. On the other hand, the model struggles with learning article-specific orientation because the annotations in the test data are partially independent of the source annotations in the training data.

6 Conclusion

We propose POLOR, a fined-tuned BERT model augmented with multiple contrastive learning objectives to deduce article-specific labels even with training data annotated based on sources. Our model shows improved performance in predicting political orientation from unseen news sources by leveraging the similarities and differences between sentences. More specifically, the model tunes sentence embeddings by contrasting them to similar/different sentences. We introduce two datasets that span several U.S. news sources to evaluate the model performance. Moreover, our annotation results exhibit the relationship between sources and opinion articles by showing that opinion news pieces are usually driven by the author's beliefs, not the source. In future work, we attempt to explore the problem by incorporating additional knowledge other than text such as the author's background, personal experiences, and social interactions.

References

- [Bahdanau, Cho, and Bengio 2014] Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. *arXiv e-prints* arXiv:1409.0473.
- [Baly et al. 2020] Baly, R.; Da San Martino, G.; Glass, J.; and Nakov, P. 2020. We Can Detect Your Bias: Predicting the Political Ideology of News Articles. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 4982–4991. Online: Association for Computational Linguistics.
- [Budak, Goel, and Rao 2016] Budak, C.; Goel, S.; and Rao, J. M. 2016. Fair and Balanced? Quantifying Media Bias through Crowdsourced Content Analysis. *Public Opinion Quarterly* 80(S1):250–271.
- [Card et al. 2015] Card, D.; Boydston, A. E.; Gross, J. H.; Resnik, P.; and Smith, N. A. 2015. The media frames corpus: Annotations of frames across issues. In *Annual Meeting of the Association for Computational Linguistics*.
- [Chen et al. 2018] Chen, W.-F.; Wachsmuth, H.; Al-Khatib, K.; and Stein, B. 2018. Learning to Flip the Bias of News Headlines. In *Proceedings of the 11th International Conference on Natural Language Generation*, 79–88. Tilburg University, The Netherlands: Association for Computational Linguistics.
- [Chen et al. 2020] Chen, W.-F.; Al Khatib, K.; Wachsmuth, H.; and Stein, B. 2020. Analyzing Political Bias and Unfairness in News Articles at Different Levels of Granularity. In *Proceedings of the Fourth Workshop on Natural Language Processing and Computational Social Science*, 149–154. Online: Association for Computational Linguistics.
- [de Jong 2023] de Jong, B. 2023. Why journalists are speaking out against western media bias in reporting on israel-palestine. Illustration by Walker Gawande, Edited by Tina Lee.
- [Fan et al. 2019] Fan, L.; White, M.; Sharma, E.; Su, R.; Choubey, P. K.; Huang, R.; and Wang, L. 2019. In plain sight: Media bias through the lens of factual reporting. 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)*, 6343–6349. Hong Kong, China: Association for Computational Linguistics.
- [Gangula, Duggenpudi, and Mamidi 2019] Gangula, R. R.; Duggenpudi, S. R.; and Mamidi, R. 2019. Detecting political bias in news articles using headline attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, 77–84. Florence, Italy: Association for Computational Linguistics.
- [Gentzkow and Shapiro 2006] Gentzkow, M., and Shapiro, J. 2006. What Drives Media Slant? Evidence from U.S. Daily Newspapers. *Econometrica* 78:35–71.
- [Greene and Resnik 2009] Greene, S., and Resnik, P. 2009. More than Words: Syntactic Packaging and Implicit Sentiment. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, 503–511. Boulder, Colorado: Association for Computational Linguistics.
- [Groseclose and Milyo 2005] Groseclose, T., and Milyo, J. 2005. A Measure of Media Bias. *The Quarterly Journal of Economics* 120(4):1191–1237.
- [Jiang et al. 2019] Jiang, Y.; Petrak, J.; Song, X.; Bontcheva, K.; and Maynard, D. 2019. Team bertha von suttner at SemEval-2019 task 4: Hyperpartisan news detection using ELMo sentence representation convolutional network. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, 840–844. Minneapolis, Minnesota, USA: Association for Computational Linguistics.
- [Kiesel et al. 2019] Kiesel, J.; Mestre, M.; Shukla, R.; Vincent, E.; Adineh, P.; Corney, D.; Stein, B.; and Potthast, M. 2019. SemEval-2019 task 4: Hyperpartisan news detection. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, 829–839. Minneapolis, Minnesota, USA: Association for Computational Linguistics.
- [Kim and Johnson 2022] Kim, M. Y., and Johnson, K. M. 2022. CLoSE: Contrastive Learning of Subframe Embeddings for Political Bias Classification of News Media. In *Proceedings of the 29th International Conference on Computational Linguistics*, 2780–2793. Gyeongju, Republic of Korea: International Committee on Computational Linguistics.
- [Kulkarni et al. 2018] Kulkarni, V.; Ye, J.; Skiena, S.; and Wang, W. Y. 2018. Multi-view Models for Political Ideology Detection of News Articles. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3518–3527. Brussels, Belgium: Association for Computational Linguistics.
- [Lazaridou et al. 2020] Lazaridou, K.; Löser, A.; Mestre, M.; and Naumann, F. 2020. Discovering biased news articles leveraging multiple human annotations. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, 1268–1277. Marseille, France: European Language Resources Association.
- [Lee et al. 2022] Lee, N.; Bang, Y.; Yu, T.; Madotto, A.; and Fung, P. 2022. Neus: Neutral multi-news summarization for mitigating framing bias.
- [Liu et al. 2022] Liu, Y.; Zhang, X. F.; Wegsman, D.; Beauchamp, N.; and Wang, L. 2022. POLITICS: Pretraining with same-story article comparison for ideology prediction and stance detection. In *Findings of the Association for Computational Linguistics: NAACL 2022*, 1354–1374. Seattle, United States: Association for Computational Linguistics.
- [Lundberg and Lee 2017] Lundberg, S. M., and Lee, S.-I. 2017. A unified approach to interpreting model predictions. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- [Reimers and Gurevych 2019] Reimers, N., and Gurevych, I. 2019. Sentence-bert: Sentence embeddings using siamese

bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

[Roy and Goldwasser 2020] Roy, S., and Goldwasser, D. 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)*, 7698–7716. Online: Association for Computational Linguistics.

[Schroff, Kalenichenko, and Philbin 2015] Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. Facenet: A unified embedding for face recognition and clustering. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 815–823.

[van der Maaten and Hinton 2008] van der Maaten, L., and Hinton, G. 2008. Visualizing data using t-sne. *Journal of Machine Learning Research* 9:2579–2605.