

Pierre Lilly¹; Brenner Little²; Arun Bagavathi²
¹Langston University, ²Oklahoma State University

Introduction

Online news media forums and social networks are the primary information system to rapidly disseminate the current news to the global population. Although social media forums have become an integral part of millions of people, they affect one's political opinion with the type and quantity of news they present. Exposure to such polarization can not only develop online disagreements and ideology segregation but may also lead to offline extreme or even violent activities. Thus, it is essential to develop quantitative methods to characterize social media polarization with the advancements in data science and machine learning.

In this work we present a pipeline to get political ideology of Tweets using a heuristic method. We experiment with multiple text representation algorithms to extract text embeddings of the tweets and we show with a supervised machine learning approach that the MLP classifier can predict the political ideology of tweets with up to **82%** accuracy. We breakdown our results collected using multiple text representation learning methods.

As our work is the first attempt to analyze political ideology of social media posts, there exists no labelled datasets, except the data that are heuristically labelled, to evaluate our methods. Thus in addition to our previous method, we also propose a promising few-shot learning method for future experiments to extract the political ideology labels of the tweets with very limited data available during the training phase. We give base results collected with our method on a limited hand-curated tweets in this work.

AllSides™ Media Bias Chart™

All ratings are based on online content only — not TV, print, or radio content. Ratings do not reflect accuracy or credibility; they reflect perspective only.



Figure 1. Media Bias Chart

Related Work

The existing computational approaches study online political polarization in multiple aspects using machine learning. For example, they identify polarized user communities on social media[1], analyze polarized topics[2], predict political leaning of news articles[3] with either supervised or unsupervised approaches. However, the existing literature focuses on analyzing political polarization at the user-level or topic-level or news-level. Characterizing the political leaning of social media posts has the potential to become a precursor for hate speech detection, fake news detection, and influence prediction on social media[4]. As an initial step, in this work we explore the capability of text representation learning and MLP classifier to predict the political ideology of social media posts.

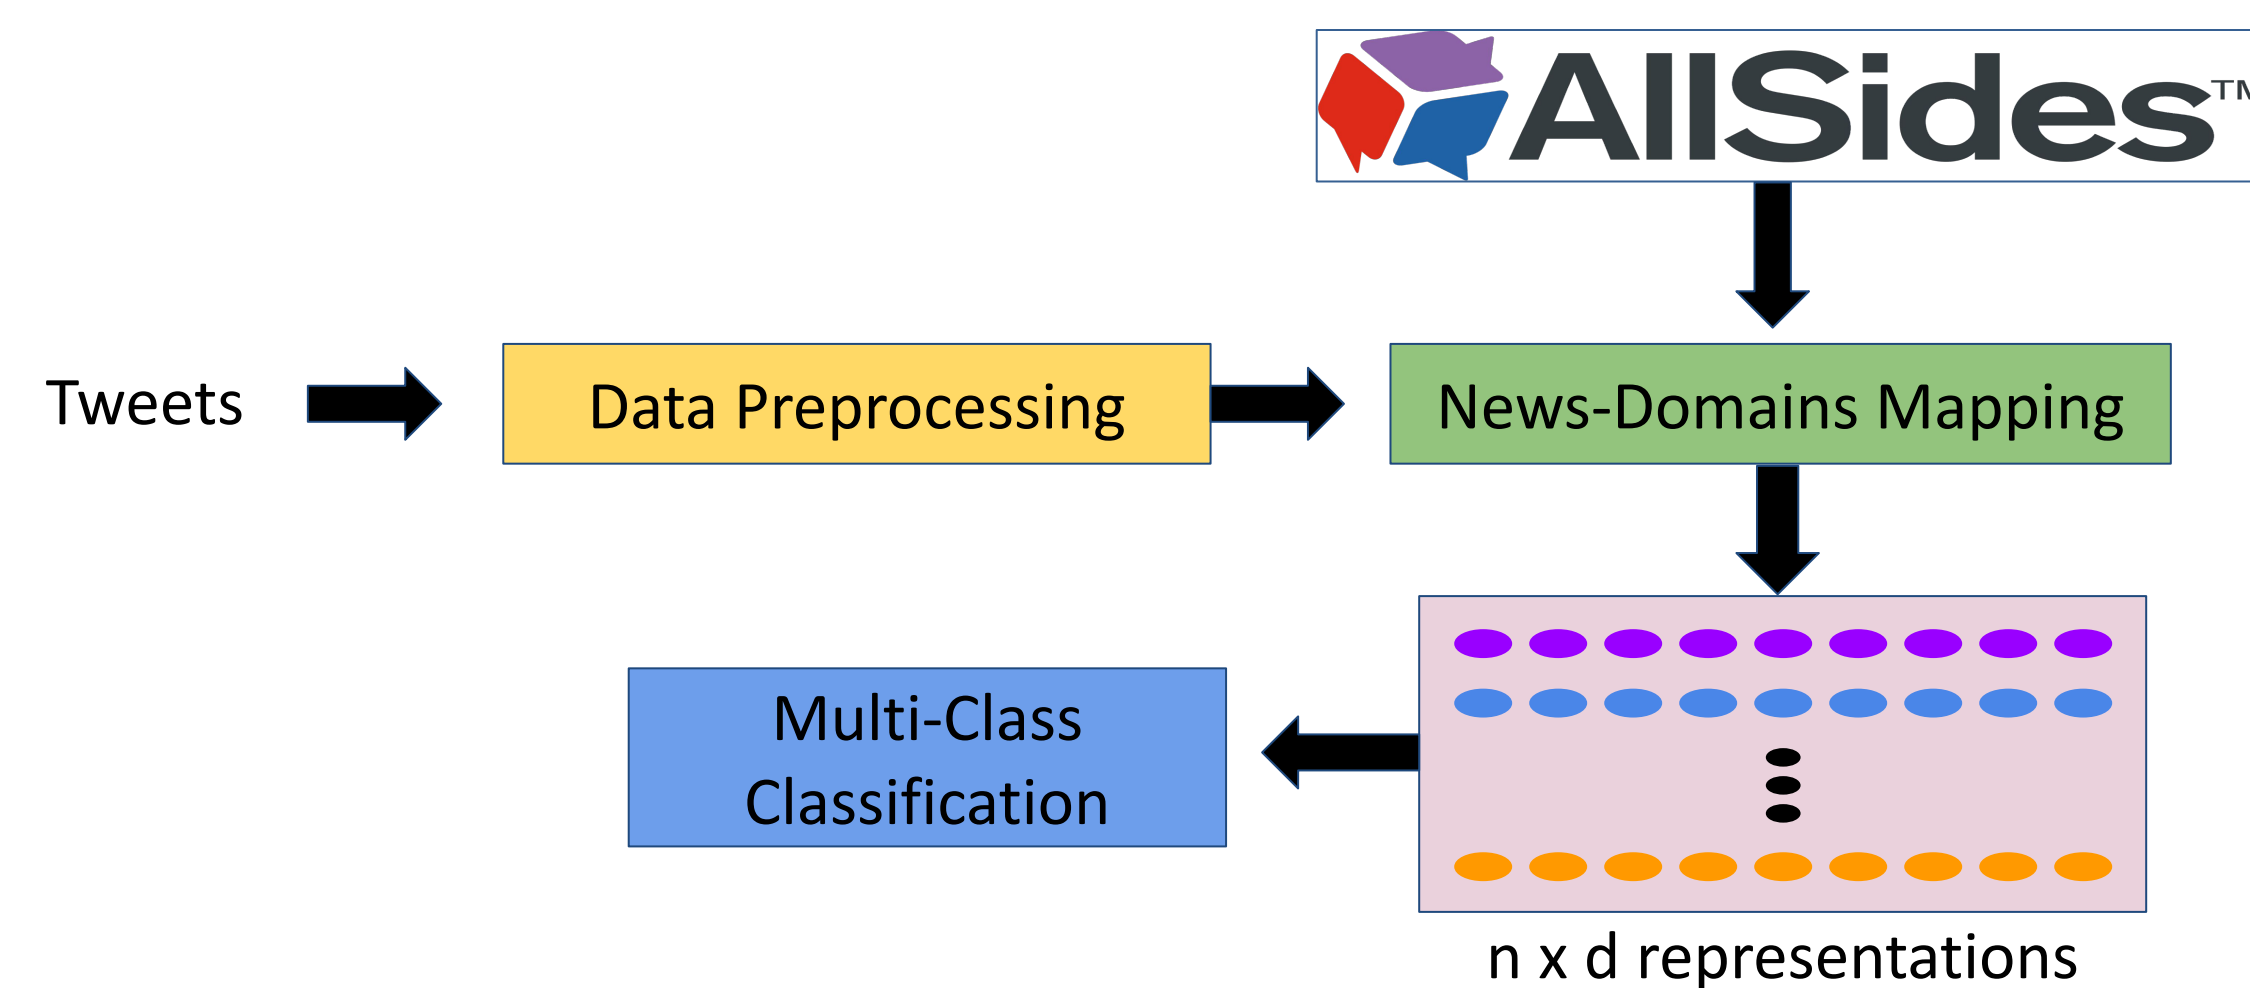


Figure 2. Methodology Pipeline

Methodology

Dataset: We use the tweets that were collected in previous research[5]. These tweets were collected focussing on the presence of news article URLs. All the datasets used in this project is summarized in *Table 1*.

Before training the model, there were data preprocessing steps done. First, all special characters and stop words were removed from the post's text. After removing those, the text was tokenized and ready to be trained in the model. The tweets are then heuristically labelled as depicted in *Figure 2*.

Each of the text embedding methods use specific hyperparameters to obtain the results. Throughout CountVectorizer, TfidfVectorizer, Word2vec, and BERT programs, the MLPClassifier is utilized with a hyperparameter for two hidden layers, 768 and 3. Additionally, the hyperparameters included a relu activation function, adam solver, alpha 1e-6, and a maximum iteration of 1000000.

The Word2Vec program had hyperparameters, including a minimum count of 1, a window of 5, and a size of 300. We used the pre-trained BERT model from HuggingFace API which gives vector representations of size 768.

Data Distribution

Dataset	Left	Center	Right
News Domains	149	166	98
Hand-curated Tweets	49	83	36

Distribution of Heuristically Labelled Tweets			
Class	Training	Test	Total
Left	167,931	41,983	209,914
Center	37,113	9,278	46,391
Right	26,746	6,687	33,433

Table 1. Datasets used in Experiments

Methods for Identifying Political Bias

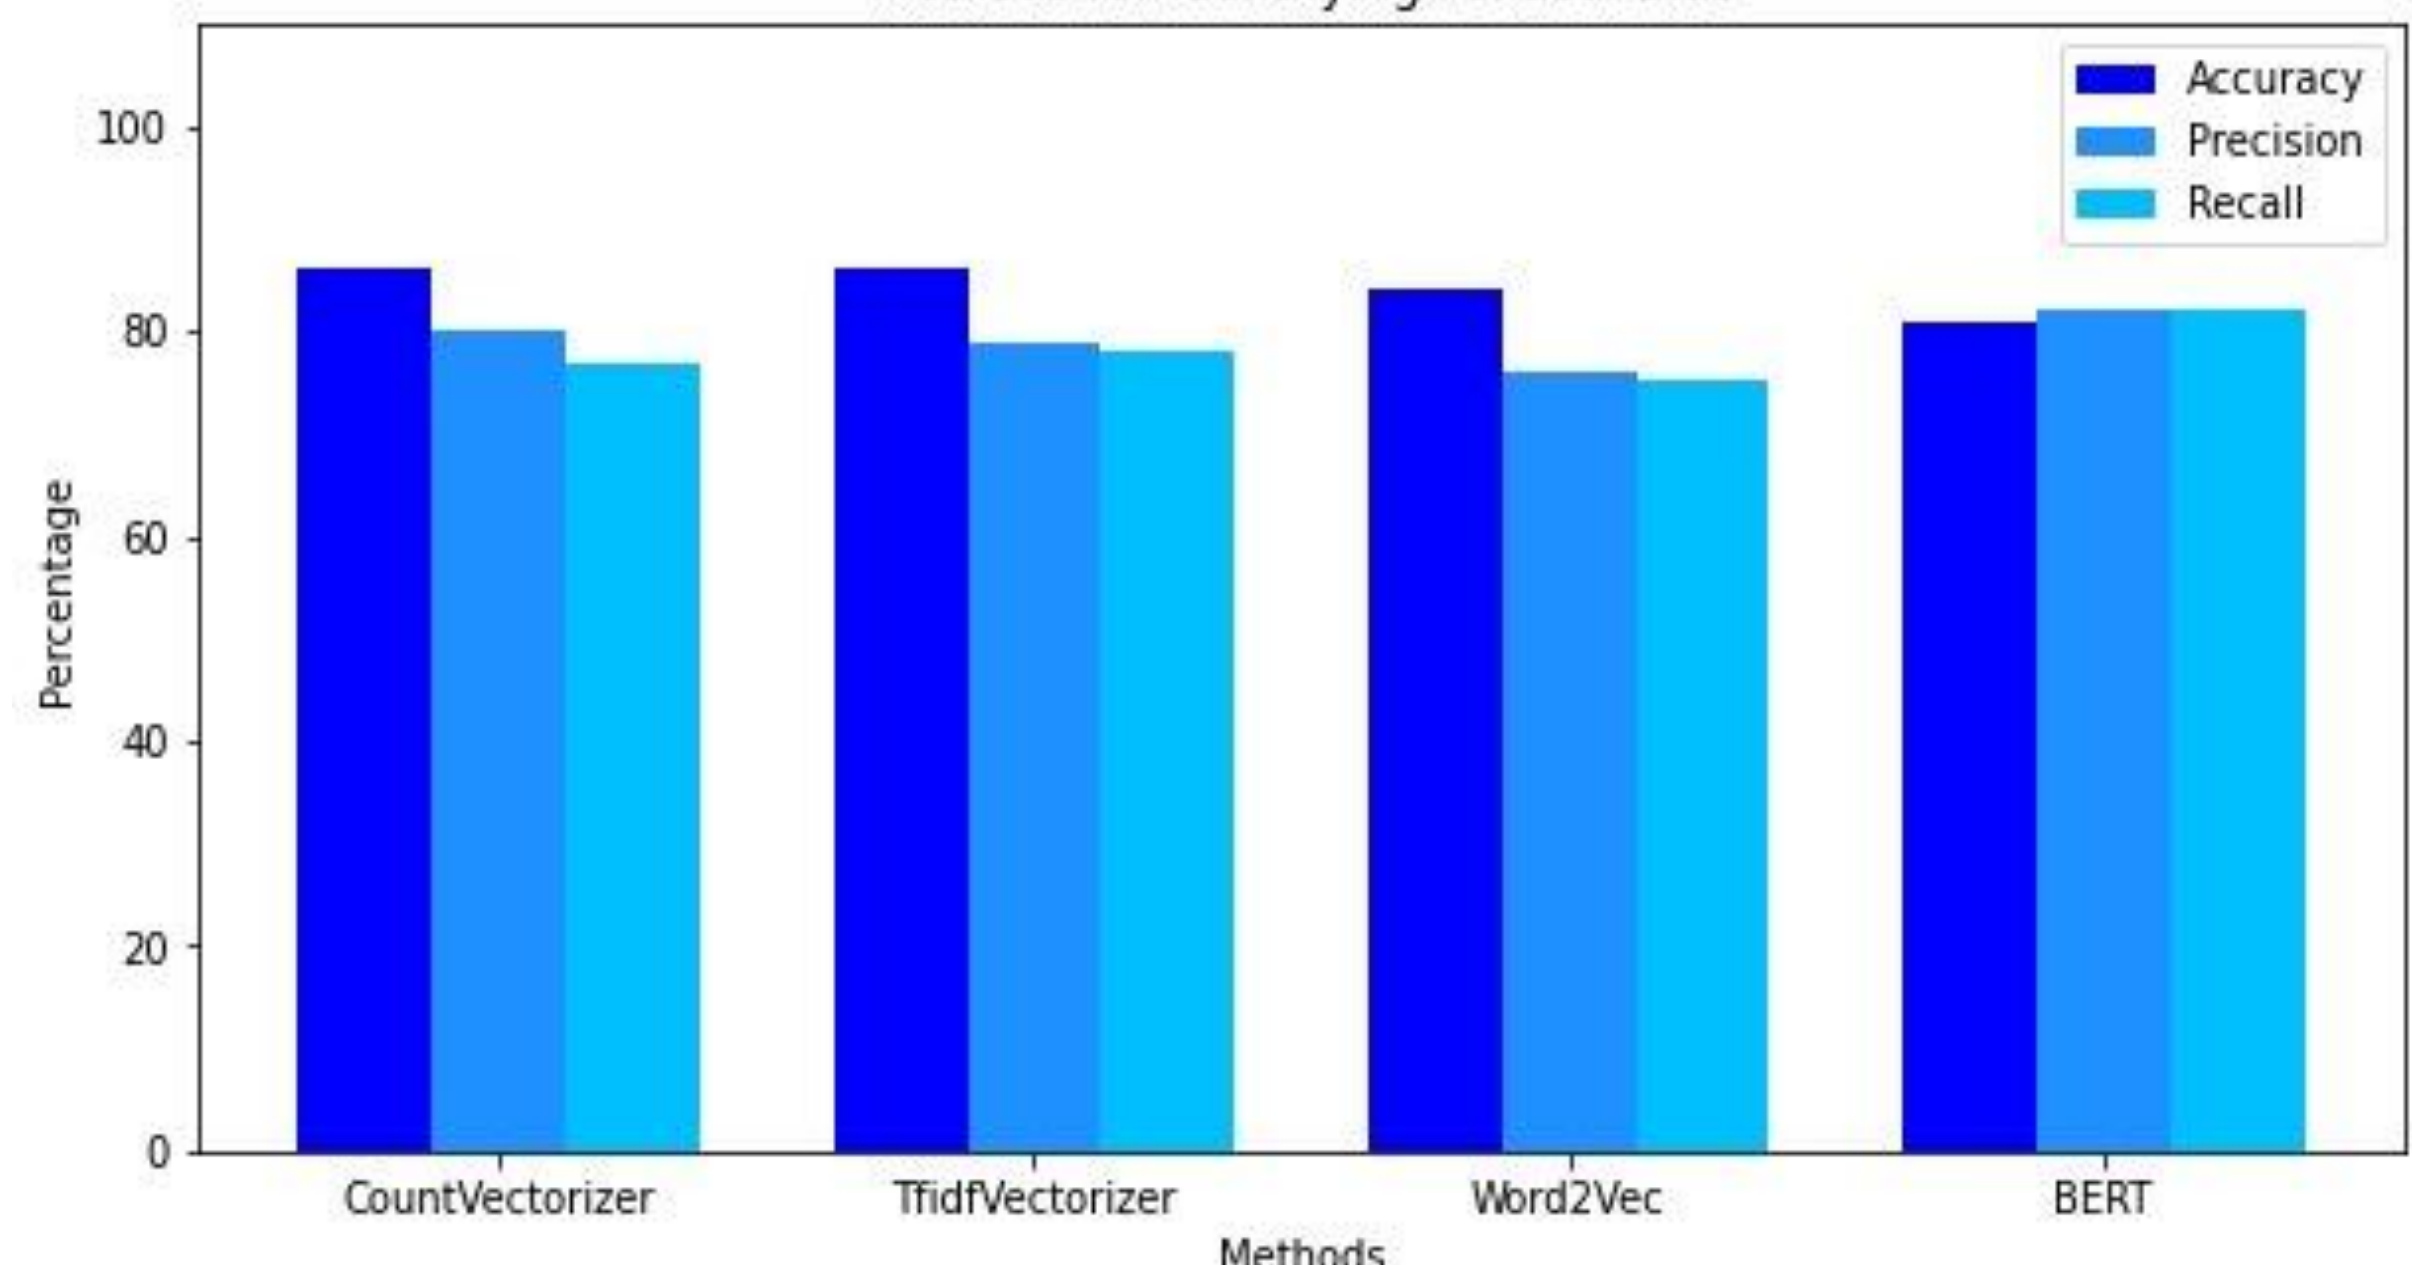


Figure 3. Bar Chart of Results

Results

We used 80% of the heuristically labelled data for training and 20% for test in all our experiments. As depicted in *Figure 3* accuracy of the MLP classifier with BERT representations is not as good as other representations. We further analyze classifier performance in terms of predicting all labels. From *Figures 4 and 5* we can notice that the MLP classifier with BERT representations can predict all labels with up to 82% accuracy.

All the above results are collected from heuristic labels. We are beginning to explore few-shot learning to obtain political ideology labels with very limited set of labeled data. Our initial results with the pre-trained BERT model on very few training set is given in *Figure 6*.



Figure 4. BERT Confusion Matrix



Figure 5. Word2Vec Confusion Matrix

Few-shot Learning Test Accuracy based on the Number of Training Samples




Figure 6. Our initial analysis on Few-Shot Learning on our hand-curated tweets

Conclusions

Overall, we show ML methods can predict the political ideology of tweets on a very large scale data. We will further explore in fine-tuning few-shot learning methods along with the domain knowledge collected from Wikipedia in upcoming weeks.

References

- Nair, Sreeja, Kin Wai Ng, Adriana Iamnitchi, and John Skvoretz. "Diffusion of social conventions across polarized communities: An empirical study." *Social Network Analysis and Mining* 11, no. 1 (2021): 1-17.
- Bagavathi, Arunkumar, Pedram Bashiri, Shannon Reid, Matthew Phillips, and Siddharth Krishnan. "Examining untempered social media: analyzing cascades of polarized conversations." In *IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pp. 625-632. 2019.
- Baly, Ramy, Giovanni Da San Martino, James Glass, and Preslav Nakov. "We can detect your bias: Predicting the political ideology of news articles." *arXiv preprint arXiv:2010.05338* (2020).
- Vicario, Michela Del, Walter Quattrociocchi, Antonio Scala, and Fabiana Zollo. "Polarization and fake news: Early warning of potential misinformation targets." *ACM TWEB* 13, no. 2 (2019): 1-22.
- Brena, Giovanni, Marco Brambilla, Stefano Ceri, Marco Di Giovanni, Francesco Pierrì, and Giorgia Ramponi. "News sharing user behaviour on twitter: A comprehensive data collection of news articles and social interactions." In *AAAI ICWSM*, vol. 13, pp. 592-597. 2019.

Contact Details

Pierre Lilly: pierre.lilly@langston.edu
Brenner Little: brenner.little@okstate.edu
Arun Bagavathi: abagava@okstate.edu