Quantifying partisanship in ideologically-biased TV news

Upasana Dutta
Department of Computer and Information Science
University of Pennsylvania, Pennsylvania, USA
upasanad@seas.upenn.edu

Homa Hosseinmardi
Department of Computer and Information Science
Annenberg School of Communication
University of Pennsylvania, Pennsylvania, USA
homahm@seas.upenn.edu

Aaron Clauset
Department of Computer Science
BioFrontiers Institute
University of Colorado Boulder
Santa Fe Institute, Santa Fe, USA
aaron.clauset@colorado.edu

Duncan Watts
Department of Computer and Information Science
Annenberg School of Communication
Operations, Information, and Decisions Department
University of Pennsylvania, Pennsylvania, USA
djwatts@seas.upenn.edu

ABSTRACT

Television news continues to be a vital information source for many Americans. Recent studies reveal an increasing trend of viewers gravitating towards ideologically-biased programming. This study introduces a novel method that leverages GPT-4 and active learning to quantify polarization in TV news programs. Our model shows a high correlation with human judgments and outperforms other large language models, such as LLaMa-7B and LLaMa-13B, as well as existing keyword-based methods. To enhance cost efficiency, we then employ active learning to strategically select samples for labeling, thereafter training a model using much fewer samples than a random sampling strategy, achieving a higher accuracy with lesser number of labeled samples. Finally, by analyzing longitudinal data from show transcripts, we quantify partisanship across all episodes from six major TV news channels in 2018, illustrating how partisanship differed between the TV channels over time.

CCS CONCEPTS

• **Computing Methodologies** → **Modeling and Simulation**; **Machine Learning**.

KEYWORDS

LLMs, active learning, partisanship, TV news

ACM Reference Format:
Upasana Dutta, Homa Hosseinmardi, Aaron Clauset, and Duncan Watts. 2018. Quantifying partisanship in ideologically-biased TV news. In *Proceedings of* . ACM, New York, NY, USA, 3 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Television news remains a significant source of information for many Americans, often surpassing online news platforms in viewership and engagement. Research indicates that ideological segregation among television audiences tends to be more prevalent, concentrated, and consistent compared to online platforms [5]. Cable news channels with distinct partisan leanings, such as Fox News and MSNBC, continue to attract and engage larger audiences over time. As the proportion of television news consumers tuning into ideologically biased programming rises, quantifying the degree of polarization in TV news broadcasts becomes essential for downstream analyses, such as studying the effects of exposure to highly partisan or ideologically slanted television news programs.

In this study, we aim to quantify partisanship in TV news using large language models, particularly GPT-4. Our analysis is conducted at the level of individual episodes, diverging from the channel-level approach of previous studies [4], since the degree of partisanship can vary considerably between different programs on the same channel. First, we demonstrate that GPT-4-based measures not only achieve a high correlation with human judgments, but also outperform other large language models and existing state-of-the-art methods [2, 4]. We then employ an active learning framework to train a multi-layer perceptron, significantly reducing annotation costs compared to labeling samples uniformly at random.

2 DATA AND METHODS

Our primary dataset encompasses US television program transcripts for the year 2018, obtained from TVEyes. We extracted transcripts of 35,449 unique news episodes from 6 major channels, three of which are broadcast networks (ABC, CBS, and NBC) and the other three are primary cable news channels (CNN, Fox News, and MSNBC). Each episode is broken down into segments of about 800 words each (i.e. approximately 5 minutes of the entire episode), with a total of 294,903 segments.

2.1 GPT-4 prompting

To measure the partisanship present in each segment, we employ GPT-4 zero-shot prompting. For each segment, we prompt GPT-4 with the following: Based on the content of the given TV show segment, provide a score for political partisanship in the transcript. The score should be on a continuous scale between 0 and 100, where

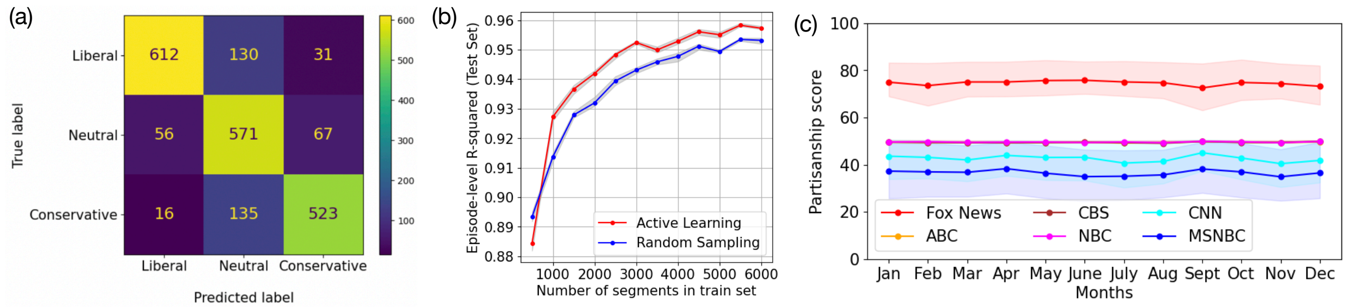

Figure 1: LLM and Active Learning Framework applied to TV news transcripts of 2018

0 means "strong bias towards liberal viewpoints" and 100 means "strong bias towards conservative viewpoints". A score of 50 means "Non-partisan viewpoints."

To validate GPT-4's performance against human annotations, we compare its responses to human labels for TV news transcripts from CNN, Fox News, and MSNBC, obtained from a previous study [3]. In this study, human annotators classified the transcripts as either "Liberal", "Neutral", or "Conservative". The findings, illustrated in Figure 1a, indicate a strong alignment between GPT-4 responses and human annotations. Both the accuracy and F1-score are 0.78, with a correlation coefficient of 0.775. These results significantly outperform other large language models, such as LLaMa-7B and LLaMa-13B, and conventional text-scaling methods that leverage Congressional speech partisanship to quantify televised news partisanship [2] (which recorded an accuracy of 0.39, and F1-score and correlation coefficient of 0.38 and 0.18, respectively).

Despite GPT-4's impressive zero-shot performance in quantifying partisanship in TV news, obtaining responses for each of the 294,903 segments would be prohibitively expensive. Therefore, we propose a novel approach of using GPT-4 to label a subset of the segments and then train a multi-layer perceptron to predict the partisanship in the remaining unlabeled segments. To minimize annotation costs efficiently, we leverage active learning to selectively label the most informative segments, thereby optimizing the performance of our multi-layer perceptron model.

2.2 Active Learning

We first compare the performance of three models—elastic net, multi-layer perceptron, and BERT—on a uniformly random subset of our dataset to determine the most suitable model for use with active learning. We randomly sample 2,000 episodes and obtain their true labels from GPT-4. This dataset is split into a training set with 1,700 episodes (14,565 segments) and a test set with 300 episodes (2,435 segments). Hyper-parameter tuning is conducted using K-fold cross-validation. The predictions on the test segments are aggregated to provide partisanship scores at the episode level.

The R^2 values for the elastic net, multi-layer perceptron, and BERT models on the test set episodes are 0.897, 0.939, and 0.834, respectively. The multi-layer thus perceptron outperformed the other models, leading us to select it for our active learning framework.

Active learning [1, 6, 7] allows us to efficiently select the most informative samples to label from a large pool of unlabeled TV news segments, enabling us to train a model with fewer samples as compared to labeling samples at random. We begin by setting aside 10% of the dataset as a test set and then employ the well-studied Query-by-Committee (QBC) framework on the remaining unlabeled samples as follows:

- We start by labeling a random sample of 500 segments.
- Iteratively, we train multiple bootstrapped models using the data labeled so far. Each model makes then predictions for every unlabeled sample. The samples with the highest disagreement (variance) between model predictions are chosen for labeling.
- This process is repeated until we exhaust our budget or achieve a predefined performance threshold.

3 RESULTS

To illustrate the advantages of active learning over random sampling for labeling data points, Figure 1(b) presents a comparative analysis of their performance. The segment-level scores are averaged to the episode level. Across all training set sizes, the Active Learning (QBC) approach consistently achieves higher R^2 compared to random sampling. Particularly, active learning crosses the 0.95 threshold with only 3000 labeled segments, compared to 4500 segments (50% more) in the random sampling case. This demonstrates that by selectively sampling the most informative segments for labeling, Active Learning leads to higher accuracy with lesser training data. Thus, in scenarios like ours where labeling data is expensive, Active Learning proves to be highly effective and resource-efficient.

Building on the promising performance of our active learning framework, we utilize the trained model from the final iteration (6000 segments) to predict the partisanship scores of the remaining unlabeled segments from 2018. Figure 1(c) presents the partisanship scores for the six TV news networks over the year. The results show that Fox News consistently exhibits the most conservative partisanship scores, reflecting a strong bias. Conversely, MSNBC and CNN display a marked liberal bias, though not to an extreme degree. Broadcast news channels (ABC, CBS, NBC), however, largely maintain a non-partisan stance.

Our project is ongoing, with plans to extend our framework to a larger dataset covering the period from 2013 to 2022 (328,432 unique episodes). The next phase of our research will involve a detailed analysis of partisanship scores at the program level, particularly focusing on ideologically-driven TV news programs that disseminate opinionated content rather than objective news reporting.

REFERENCES

[1] Mitchell Bosley, Saki Kuzushima, Ted Enamorado, and Yuki Shiraito. 2022. Improving Probabilistic Models in Text Classification via Active Learning. *arXiv preprint arXiv:2202.02629* (2022).

[2] Matthew Gentzkow and Jesse M Shapiro. 2010. What drives media slant? Evidence from US daily newspapers. *Econometrica* 78, 1 (2010), 35–71.

[3] Eunji Kim, Yphtach Lelkes, and Joshua McCrain. 2022. Measuring dynamic media bias. *Proceedings of the National Academy of Sciences* 119, 32 (2022), e2202197119.

[4] Gregory J Martin and Ali Yurukoglu. 2017. Bias in cable news: Persuasion and polarization. *American Economic Review* 107, 9 (2017), 2565–2599.

[5] Daniel Muise, Homa Hosseinmardi, Baird Howland, Markus Mobius, David Rothschild, and Duncan J Watts. 2022. Quantifying partisan news diets in Web and TV audiences. *Science advances* 8, 28 (2022), eabn0083.

[6] Keze Wang, Dongyu Zhang, Ya Li, Ruimao Zhang, and Liang Lin. 2016. Cost-effective active learning for deep image classification. *IEEE Transactions on Circuits and Systems for Video Technology* 27, 12 (2016), 2591–2600.

[7] Lin Yang, Yizhe Zhang, Jianxu Chen, Siyuan Zhang, and Danny Z Chen. 2017. Suggestive annotation: A deep active learning framework for biomedical image segmentation. In *Medical Image Computing and Computer Assisted Intervention-MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11-13, 2017, Proceedings, Part III 20*. Springer, 399–407.