

Political Ideology and Polarization: A Multi-dimensional Approach

Barea Sinno^{1*} Bernardo Oviedo^{2*} Katherine Atwell^{3*}

Malihe Alikhani³ Junyi Jessy Li⁴

¹ Political Science, Rutgers University

² Computer Science, ⁴ Linguistics, The University of Texas at Austin

³ Computer Science, University of Pittsburgh

barea.sinno@gmail.com, bernyoviedo@utexas.edu, kaa139@pitt.edu
malihe@pitt.edu, jessy@utexas.edu

Abstract

Analyzing ideology and polarization is of critical importance in advancing our grasp of modern politics. Recent research has made great strides towards understanding the ideological *bias* (i.e., stance) of news media along the left-right spectrum. In this work, we instead take a novel and more nuanced approach for the study of ideology based on its left or right positions on the issue being discussed. Aligned with the theoretical accounts in political science, we treat ideology as a multi-dimensional construct, and introduce the first diachronic dataset of news articles whose ideological positions are annotated by trained political scientists and linguists at the paragraph level. We showcase that, by controlling for the author’s stance, our method allows for the quantitative and temporal measurement and analysis of polarization as a multidimensional ideological distance. We further present baseline models for ideology prediction, outlining a challenging task distinct from stance detection.

1 Introduction

Political ideology rests on a set of beliefs about the proper order of a society and ways to achieve this order (Jost et al., 2009; Adorno et al., 2019; Campbell et al., 1980). In Western politics, these worldviews translate into a multi-dimensional construct that includes: equal opportunity as opposed to economic individualism; general respect for tradition, hierarchy and stability as opposed to advocating for social change; and a belief in the un/fairness and in/efficiency of markets (Jost et al., 2009).

The divergence in ideology, i.e., polarization, is the undercurrent of propaganda and misinformation (Vicario et al., 2019; Bessi et al., 2016; Stanley, 2015). It can congest essential democratic functions with an increase in the divergence of political ideologies. Defined as a growing ideo-

Two dimensions: trade and economic liberalism	The U.S. aim is to create a monetary system with enough flexibility to prevent bargain-hungry money from rolling around the world like loose ballast on a ship disrupting normal trade and currency flows. Nixon goals: dollar, trade stability. This must be accompanied, Washington says, by reduction of [trade] barriers ...
One dimension: trade protectionism	The controls program, which Mr. Nixon inaugurated Aug. 15, 1971, has helped to reduce inflation to about 3 percent yearly, and to boost annual U.S. economic growth to more than 7 percent...

Table 1: Excerpts from news article #730567 in COHA (Davies, 2012). The first paragraph advocates for liberalism and the reduction of trade barriers. It also has a domestic economic dimension. The second paragraph, on the contrary, advocates for protectionism and a domestic controls program.

logical distance between groups, polarization has waxed and waned since the advent of the American Republic (Pierson and Schickler, 2020).¹ Two eras—post-1896 and -1990s—have witnessed deleterious degrees of polarization (Jenkins et al., 2004; Jensen et al., 2012). More recently, COVID-19, the murder of George Floyd, and the Capitol riots have exposed ideological divergences in opinion in the US through news media and social media. With the hope of advancing our grasp of modern politics, we study ideology and polarization through the lens of computational linguistics by presenting a carefully annotated corpus and examining the efficacy of a set of computational and statistical analyses.

In contrast to studying the bias or the stance of the author of the text via linguistic framing (Kulkarni et al., 2018; Kiesel et al., 2019; Baly et al., 2019, 2020; Chen et al., 2020; Stefanov et al., 2020), we

¹We distinguish ourselves from work that considers other types of polarization, e.g., as a measure of emotional distance (Iyengar et al., 2019) or distance between political parties (Lauka et al., 2018).

* Equal contribution ordered by first name.

study the little explored angle that is nonetheless critical in political science research: **ideology of the issue (e.g., policy or concept) under discussion.** That is, in lieu of examining the author’s stance, we focus on addressing the at-issue content of the text and the ideology that it represents in the implicit social context. The nuanced co-existence of stance and ideology can be illustrated in the following excerpt:

“Republicans and Joe Biden are making a huge mistake by focusing on cost. The implication is that government-run health care would be a good thing—a wonderful thing!—if only we could afford it.” (The Federalist, 9/27/2019)

The author is attacking a liberal social and economic policy; therefore, the ideology being discussed is liberal on two dimensions—social and economic, while the author’s stance is conservative. Moreover, our novel approach acknowledges that ideology can also *vary* within one article. In Table 1, we show an example in which one part of an article advocates for trade liberalism, while another advocates for protectionism.

Together, author stance and ideology inform us not only that there is bias in the media, but also which beliefs are being supported and/or attacked. A full analysis of *polarization* (that reflects a growing distance of political ideology over time) can then be derived if *diachronic* data for both author stance and ideology were available. However, while there has been data for the former (with articles from recent years only) (Kiesel et al., 2019), to date, there has been no temporal data on the latter.

In this paper, we present a multi-dimensional framework, and an annotated, diachronic, stance-neutral corpus, for the analysis of ideology in text. This allows us to study polarization as a *state* of ideological groups with divergent positions on a political issue as well as polarization as a *process* whose magnitude grows over time (DiMaggio et al., 1996). We use proclaimed center, center-left and center-right media outlets who claim to be objective in order to focus exclusively and more objectively on the ideology of the issue being discussed, without the subjectivity of author stance annotation. We study ideology within every paragraph² of an article and aim to answer the following question: which ideological dimension is present and to which ideological position does it correspond to on the liberal-conservative spectrum.

²We use automatically segmented paragraphs since the raw texts were not paragraph-segmented.

Our extensive annotation manual is developed by a political scientist, and the data then annotated by three linguists after an elaborate training phase (Section 3). After 150 hours of annotation, we present a dataset of 721 *fully adjudicated* annotated paragraphs, from 175 news articles and covering an average of 7.86 articles per year (excerpts shown in Tables 1, 2, and 3). These articles originate from 5 news outlets related the US Federal Budget from 1947-1975 covering the center-left, center, center-right spectrum: Chicago Tribune (CT), Christian Science Monitor (CSM), the New York Times (NYT), Time Magazine (TM), and the Wall Street Journal (WSJ).

With this data, we reveal lexical insights on the language of ideology across the left-right spectrum and across dimensions. We observe that linguistic use even at word level can reveal the ideology behind liberal and conservative policies (Section 4). Our framework also enables fine-grained, quantitative analysis of polarization, which we demonstrate in Section 5. This type of analysis, if scaled up using accurate models for ideology prediction, has the potential to reveal impactful insights into the political context of our society as a whole.

Finally, we present baselines for the automatic identification of multi-dimensional ideology at the paragraph level (Section 6). We show that this is a challenging task with our best baseline yielding an F measure of 0.55; exploring pre-training with existing data in news ideology/bias identification, we found that this task is distinct from, although correlated with, labels automatically derived from news outlets. *We contribute our data and code at <https://github.com/bernovie/political-polarization>.*

2 Setup

Many political scientists and political psychologists argue for the use of *at least* a bidimensional ideology for domestic politics that distinguishes between economic and social preferences (Carsey and Layman, 2006; Carmines et al., 2012; Feldman and Johnston, 2014).³ We start with these

³It is important to distinguish between ideology and several other concepts. (1) **Partisanship** (party identity) (Campbell et al., 1980): a partisan person changes their ideology when their party changes its ideology, whereas an ideological person changes their party when their party changes its ideology. Partisanship is easily conflated with party ID using a unidimensional conceptualization of ideology, but not with a multi-dimensional one. (2) **Moral foundations**: Haidt et al. (2009) gave an evolutionary explanation of how human

two dimensions while adding a third dimension, “Foreign”, when the article tackles foreign issues.

Specifically, our **annotation task** entails examining a news article and annotating each dimension (detailed below) along three levels—*liberal, conservative, neutral*—for each paragraph. The neutral level for every dimension is reserved for paragraphs related to a specific dimension but either (a) contain both conservative and liberal elements that annotators were unable to ascertain an ideological dimension with confidence, or (b) do not portray any ideology. We additionally provide an *irrelevant* option if a dimension does not apply to the paragraph. The three dimensions are:

Social: While the (1) *socially conservative* aspect of this dimension is defined as respect for tradition, fear of threat and uncertainty, need for order and structure, concerns for personal and national security, and preference for conformity, its (2) *socially liberal* counterpart has been associated with a belief in the separation of church and state, tolerance for uncertainty and change (Jost et al., 2009).

Economic: Similarly, while the (3) *economically conservative* aspect of this dimension refers to motivations to achieve social rewards, power, and prestige such as deregulation of the economy, lower taxes and privatization (i.e., being against deficit) spending and advocating for a balanced budget, its (4) *economically liberal* counterpart refers to motivation for social justice and equality such as issues related to higher taxes on rich individuals and businesses and more redistribution.

Foreign: After piloting the bidimensional approach on 300 articles, we find that using only 2 dimensions conflates two important aspects of ideology related to domestic economy and foreign trade. Tariffs, import quotas, and other nontariff-based barriers to trade that are aimed at improving employment and the competitiveness of the US on the international market did not map well onto the bidimensional framework. After consulting several senior political scientists, we adopted a third dimension that dealt with the markets as well as the relations of the US with the rest of the world. While

morals, values and traits such as freedom, safety, harm, care, reciprocity, in-group loyalty, authority, equality are formed. Since, some scholars have used these traits to predict ideology whereas others have attempted to understand what traits unites people with the same ideology. (3) **Framing:** frames are used in many ways in political science. They can refer to different ways scholars describe the same information or when scholars talk about different aspects of a single problem (Chong and Druckman, 2007).

the (5) *globalist* counterpart of this dimension accounts for free-trade, diplomacy, immigration and treaties such as the non-proliferation of arms, its (6) *interventionist* aspect is nationalist in its support for excise tax on imports to protect American jobs and economic subsidies and anti-immigration.

With the annotated data, we demonstrate quantitative **measures of polarization** (Section 5) and introduce the **modeling task** (Section 6) of automatically identifying the ideology of the policy positions being discussed.

3 Data collection and annotation

Raw data Since polarization is a *process* that needs to be analyzed *over time* (DiMaggio et al., 1996), our annotated articles are sampled from a diachronic corpus of 1,749 news articles across nearly 3 decades (from 1947 till 1974). Articles in this corpus are from political news articles of Desai et al. (2019) from the Corpus of Historical American English (COHA, Davies (2012)) covering years 1922-1986. These 1,749 articles are extracted such that: (1) they cover broad and politically relevant topics (ranging from education and health to economy) but still share discussions related to the federal budget to make our annotations tractable⁴; (2) balanced in the number of articles across 5 news outlets with center-left, central, and center-right ideology (c.f. Section 5): Chicago Tribune (CT), Wall Street Journal (WSJ), Christian Science Monitor (CSM), the New York Times (NYT), and Time Magazine (TM). A detailed description of our curation process is in Appendix A.

The raw texts were not segmented into paragraphs, thus we used Topic Tiling (Riedl and Biemann, 2012) for automatic segmentation. Topic Tiling finds segment boundaries again using LDA and, thus, identifies major subtopic changes within the same article. The segmentation resulted in articles with 1 to 6 paragraphs. The average number of paragraphs per article was 4.

Annotation process Our team (including a political science graduate student) developed an annotation protocol for expert annotators using definitions in Section 2. The annotation process is independently reviewed by four political science professors from two universities in the US who are not authors

⁴Because federal budget stories touch on all aspects of the federal activity, this topic appeals to both liberal and conservative media and thus can provide a good testing ground to showcase our proposed ideological annotation method.

Two dimensions: socially and economically liberal	... Secretary of Defense Robert S. [McNamara] threw his full support today behind the Administration's drive against poverty. ...Mr. Mc-Namara said : "It is the youth that we can expect to be the most immediate beneficiaries of the war on poverty." He said he was endorsing the "entire program" both as a citizen and as a member of the Cabinet. His endorsement came as his fellow Republicans in Congress continued to hammer away at parts of the Administration's antipoverty program. ...	Zero dimension	... "The committee is holding public hearings on President Eisenhower's Economic Report, which he sent to Congress last week. The Secretary's [Humphrey] appearance before the group provided an opportunity for political exchanges.
Two dimensions: socially and economically conservative	The antipoverty program, the Republicans insisted, would undercut the authority of the Cabinet members by making Sargent Shriver a "poverty czar." "I don't see how you can lie down and be a doormat for this kind of operation. "...	One dimension: economically liberal	Senators Paul H. Douglas of Illinois, J. W. Fulbright of Representative Wright Patman of Texas, all Democrats, were active in questioning Mr. Humphrey. The Democrats asserted that the Administration's tax reduction program was loaded in favor of business enterprises and shareholders in industry and against the taxpayer in the lowest income brackets....
		One dimension: economically neutral	Senator Fulbright .. declare[d] that the problem was to expand consumption rather than production. ... "Production is the goose that lays the golden egg," Mr. Humphrey replied. "Payrolls make consumers."

Table 2: Excerpts from article #723847 in COHA. Because the first paragraph calls for minimizing income inequality, it is socially liberal; and because advocating for such a program call for an budgetary expenditure, it is also has an economic liberal dimension. The second paragraph advocates for the exact opposites of the positions in the first paragraph. Therefore, it is socially and economically conservative. Sentences most relevant to these labels are highlighted.

of this paper; the research area of two of them is ideology and polarization in the US. We will release our full annotation interface, protocol, and procedure along with the data upon publication.

We sampled on average 7.86 articles per year for annotation, for a total of 721 paragraphs across 175 articles. We divided the annotation task into two batches of 45 and 130 articles, the smaller batch was for training purposes.

In addition to the political science graduate student, we recruited three annotators, all of whom are recent Linguistics graduates in the US. The training sessions consisted of one general meeting (all annotators met) and six different one-on-one meetings (each annotator met with another annotator once). During initial training, the annotators were asked to highlight sentences based on which the annotation was performed.

After the annotations of this batch were finalized, the annotators met with the political science student to create ground truth labels in cases of disagreement. Then, the three annotators received the second batch and each article was annotated by 2 annotators. This annotation was composed of two stages to account for possible subjectivity. In stage 1, each annotator worked on a batch that

Table 3: Excerpts from article #716033 in COHA. The first paragraph is void of ideology. In the second paragraph the topic is anti tax reduction on businesses, thus it is economically liberal. The third paragraph is simultaneously economically conservative and liberal because one speaker is advocating for decreasing tax on businesses and asserting that production gives an advantage to businesses, the other is advocating for decreasing tax on the poor because they need the income and asserting that healthy businesses are the ones who pay salaries for the low income bracket worker.

overlapped with only one other annotator. In stage 2, the two annotators examine paragraphs that they disagree, and met with the third annotator acting as consolidator to adjudicate. Tables 1, 3 and 2 are examples of adjudicated annotation in the data.

Agreement To assess the inter-annotator agreement of stage 1, we report Krippendorff's α (Hayes and Krippendorff, 2007) for each dimension for the 135 articles after training and before any discussion/adjudication: economic (0.44), foreign (0.68), social (0.39). The agreements among annotators for the economic and foreign dimensions are moderate & substantial (Artstein and Poesio, 2008), respectively; for social, the 'fair' agreement was noticed during annotation, and additional discussion for each paragraph was then held. Afterwards, 25 more articles were independently annotated and assessed with an α of 0.53. Although the agreements were not perfect and reflected a degree of subjectivity in this task, all dimensional labels were adjudicated after discussions between annotators. In total, creating this dataset cost \sim 150 hours of manual multi-dimensional labeling.

	#Docs	Econ.	Soc.	Fgn.	Total
CSM	37	115	63	82	260
CT	14	48	33	16	97
NYT	60	219	114	130	463
TM	52	134	60	89	283
WSJ	12	42	21	21	84
Total	175	558	291	338	1187

Table 4: Dimensional label counts across all 721 paragraphs in the adjudicated data (there can be multiple dimensions per paragraph).

Qualitative analysis of text highlights For the 25 articles used in training, all annotators highlighted the sentences that are relevant to each dimension they annotated. This helped annotators to focus on the sentences that drove their decision, and provided insights to the language of ideology, which we discuss here. On average, 21%–54% of the sentences in a paragraph were highlighted.

We found entities such as “President” and “Congress” were the most prevalent in the highlights, and they tackled social and economic issues combined. This is not surprising as it suggests that when the media quotes or discusses the “President” and “Congress”, they do so with reference to more complex policy issues. In contrast, individual congresspeople tackled mostly economic or social issues. This also is not surprising as it suggests that individual congresspeople are more concerned with specific issues. Interestingly, “House” and “Senate” almost always figured more in social issues. This suggests that when news media speaks about a specific chamber, they do so associating this chamber with social issues. Finally, party affiliation was infrequent and was mostly associated with social issues.

4 Ideology analyses

The number of paragraphs per dimension in total is: **Economics (558)**, **Social (291)**, **Foreign (338)**, across the 175 articles. In Table 4 we tabulate this for each of the news outlets. Figure 1 shows the dimensional label distributions per outlet for each dimension. Expectedly, the dimensional labels often diverge from proclaimed ideology of the news outlet.

We also analyze the percentage of articles that contain at least one pair of paragraph labels that lean in different directions; for instance, a paragraph with a label of globalist (i.e., liberal) in the foreign dimension and another paragraph with a

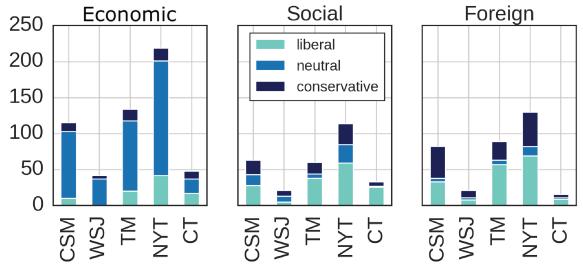


Figure 1: Dimensional label distribution per outlet.

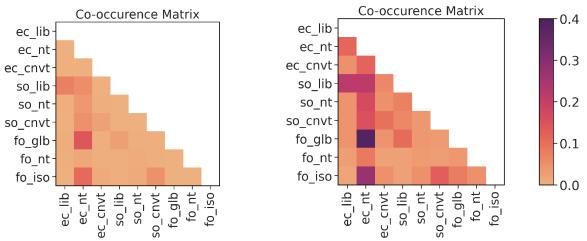


Figure 2: Co-occurrence matrices on the paragraph level (left) and article level (right)

label of conservative or neutral in the fiscal dimension. The percentage of such articles is 78.3%. Out of these articles, we examine the average proportion of neutral, liberal, and conservative paragraph labels, and find neutral labels have the highest share (43.27%), followed by liberal (33.20%) and conservative (23.53%). In Figure 2 (right), we visualize the percentage of articles where two dimensional labels co-occur within the same article. The figure indicates that ideology varies frequently within an article, showing that a single article-level label will not be fine-grained enough to capture variances within an article.

In Figure 2 (left) we also show paragraph-level label co-occurrence. Unlike the article-level, the co-occurrences are less frequent and we are more likely to observe co-occurrences along the same side of ideology. Still, we see interesting nuances; for example, on both the paragraph and the article level, the economic dimension is often neutral, and this tends to co-occur with both liberal and conservative positions in other dimensions.

Lexical analysis To understand ways ideology is reflected in text, we also look into the top vocabulary that associates with conservative or liberal ideology. To do so, we train a logistic regression model for each dimension to predict whether a paragraph is labeled conservative or liberal on that dimension, using unigram frequency as features. In Table 5 we show the top most left-leaning (L)

Economic	C: mr (5.2), tax (5.0), truman (3.7), business (3.3), billion (3.2) L: school (-4.3), education (-3.3), commission (-3.2), senator (-3.0), plan (-2.5)
Social	C: defense (7.5), tax (4.6), air (4.4), billion (3.9), missile (3.8) L: federal (-3.6), wage (-3.6), would (-3.5), policy (-3.2), labor (-2.9)
Foreign	C: defense (6.9), force (5.3), north (5.2), air (5.0), vietnam (4.6) L: aid (-9.3), economic (-5.5), foreign (-5.3), germany (-4.3), make (-4.1)

Table 5: Words with the most positive and negative weights from a logistic regression model trained to predict liberal/conservative ideology for each dimension.

or right-leaning (R) vocabulary with their weights. The table intimately reproduces our annotation of ideology. For example, words like federal and Senator allude to the fact that the topic is at the federal level. The importance of education and labor to liberals is also evident in the economic and social dimensions in words like school, education, and wage. The importance of the topic of taxation and defense is evident in conservative ideology in words such as tax, business, missile, and force.

5 Polarization

In this section, we demonstrate how our framework can be used to analyze ideological polarization, quantitatively. To say that two groups are polarized is to say that they are moving to opposite ends of an issue on the multi-dimensional ideological spectrum while, at the same time, their respective political views on ideological issues converge within a group, i.e. socially liberals become also economically liberal (Fiorina and Abrams, 2008). In political science when ideology is multidimensional, polarization is often quantified by considering three measures that capture complementary aspects (Lelkes, 2016): (1) **sorting** (Abramowitz and Saunders, 1998) (the extent to which the annotated ideology deviates from an outlet’s proclaimed ideological bias); (2) **issue constraint** (Baldassarri and Gelman, 2008) (a correlational analysis between pairs of ideological dimensions); (3) **ideological divergence** (Fiorina and Abrams, 2008) (the magnitude of the distance between two groups along a single dimension). Together these measures describe changes in the ideological environment over time: a *concurrent* increase in all three measures indicates polarization in media.

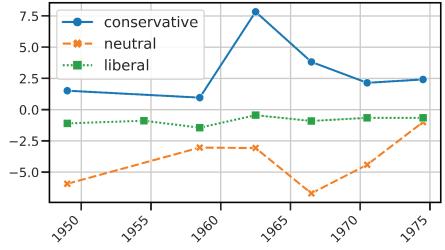


Figure 3: The evolution of the sorting measure, aggregating conservative/neutral/liberal outlets. Moving further away from the zero means articles deviate more from the proclaimed ideology of their outlets.

Limitations: We use only the fully adjudicated data and refrain from using model predictions, since our baseline experiments in Section 6 show that predicting ideology is challenging. Hence, the analysis are demonstrations of what our framework *enables*, which we discuss at the end of this section, and the conclusions are drawn for our annotated articles only. We group our data in four-year periods to reduce sparsity.

Measure 1: Sorting We adapt the sorting principle of Abramowitz and Saunders (1998) to our data and investigate the difference between the proclaimed ideological bias of a news outlet and the ideology of annotated articles from the outlet. To obtain the bias B_j of a news outlet j , we average the ratings of each news outlet across common sites that rates media bias (Adfontes, Allsides, and MBFC), yielding: **CSM (-0.07)**, **CT (0.15)**, **NYT (-0.36)**, **TM (-0.4)**, **WSJ (0.32)** (c.f. Table 8 in Appendix B for ratings from each site).

To obtain the overall ideology $I_i^{(j)}$ of article i from outlet j , we take the average of liberal (-1), neutral (0), and conservative (1) labels across its paragraphs in all three dimensions. Thus, for each 4-year time period with m articles for outlet j , the sorting measure would be the absolute distance of article vs. outlet ideology $|\text{avg}_{i=1}^m(I_i^{(j)}) - B_j|/B_j$.

In Figure 3, we plot the sorting measure, averaged across news outlets of the same proclaimed ideological bias. The figure shows that in our sample of articles, the left-leaning news outlets were closest to their proclaimed ideological bias measure over time, whereas the neutral outlets were more liberal before 1957 and after 1964. The right-leaning outlets were more conservative at that time than their proclaimed ideological bias.

Measure 2: Issue constraint This measure refers to the tightness between ideological dimen-

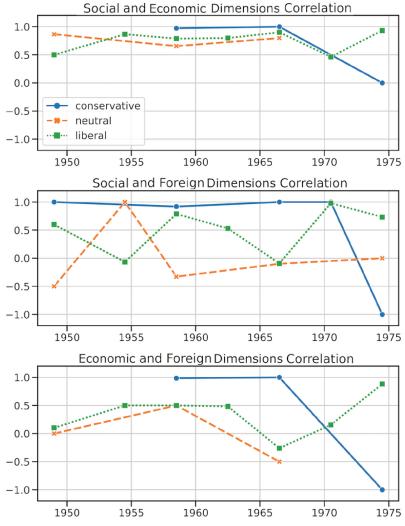


Figure 4: The evolution of the issue constraint measure, stratified by pairs of dimensions. Higher values mean some dimensions correlate more strongly than others. Due to the lack of articles that simultaneously contains social & economic dimensions (1st graph), and economic & foreign dimensions (3rd graph) from conservative outlets their respective blue lines start in 1958.

sions over time (Baldassari and Gelman, 2008) so as to assess, for example, if socially liberal dimensions are more and more associated with economic liberal dimensions for the news outlets. Concretely, for each article we derive its ideology along a single dimension as the average of paragraph annotations along that dimension. We, then, calculate the Pearson correlation between the article ideology of each pair of dimensions, over all articles from one outlet in the same period.

Results in Figure 4, again averaged across news outlets of the same ideological bias, show that for right-leaning media, the correlation between any two dimensions in the annotated data are largely positive (e.g., economically conservative were also socially conservative) until 1967 or 1970. However, for proclaimed left-leaning and neutral outlets, the correlations fluctuates especially when considering the foreign dimension.

Measure 3: Ideological divergence This measures the distance between two ideological groups on a single dimension (Fiorina and Abrams, 2008). We follow Lelkes (2016) and calculate the bimodality coefficient (Freeman and Dale, 2013; Pfister et al., 2013) per dimension over articles from all news outlets over the same time period. The bimodality coefficient ranges from 0 (unimodal, thus not at all polarized) to 1 (bimodal, thus completely

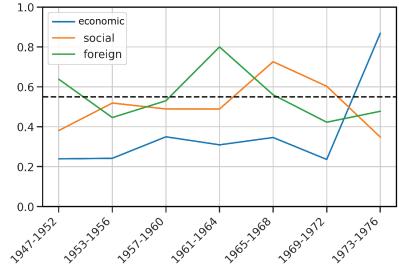


Figure 5: The evolution of the ideological divergence measure stratified by dimension. The dotted line refers to the bimodality threshold (Lelkes, 2016). Higher values mean the ideology of an article along that one dimension is bimodal.

polarized).

Figure 5 shows the evolution of the ideological divergence measure of every dimension. A bimodality measure assesses whether this divergence attained the threshold for the cumulative distribution to be considered bimodal. Ideological distance, as a result, refers to the three bimodal coefficients. We note, for example, that the foreign dimension crossed this threshold between 1956 and 1968. This means that proclaimed left-leaning and right-leaning outlets grew further apart on foreign issues during this time period.

Discussion Taken together, the graphs indicate that the years between 1957 and 1967 are the most noteworthy. During this period, from our sample of articles, we see that polarization was only present in conservative news media because it (1) sorted, as it was significantly more conservative than its composite bias measure, (2) constrained its issues, as evidence by high positive correlation values, and (3) became increasingly bimodal, as the ideological distance between their positions and those of their liberal counterpart on foreign issues increased over time. While this conclusion applies to only the set of articles in our dataset, the above analysis illustrates that our framework enables nuanced, quantitative analyses into polarization. We leave for future work, potentially equipped with strong models for ideology prediction, to analyze the data at scale.

6 Experiments

We present political ideology detection experiments as classification tasks per-dimension on the paragraph level.

We performed an 80/10/10 split to create the train, development, and test sets. The development

and test sets contain articles uniformly distributed from our time period (1947 to 1974) such that no particular decade is predominant. To ensure the integrity of the modeling task, all paragraphs belonging to the same article are present in a single split. The number of examples in the splits for each dimension for the adjudicated data are as follows: for the economic dimension, we had 450 training, 50 development, and 58 test examples. For the social dimension, we had 253 for training, 13 for development, and 25 for testing. For the foreign dimension, we had 266 for training, 33 for development, and 39 for testing.

6.1 Models

Recurrent neural networks We trained a 2-layer bidirectional LSTM (Hochreiter and Schmidhuber, 1997), with sequence length and hidden size of 256, and 100D GloVe embeddings (Pennington et al., 2014).

Pre-trained language models We used BERT-base (Devlin et al., 2019) from HuggingFace (Wolf et al., 2020) and trained two versions, with and without fine-tuning. In both cases we used a custom classification head consisting of 2 linear layers with a hidden size of 768 and a ReLU between them. To extract the word embeddings we followed Devlin et al. (2019) and used the hidden states from the second to last layer. To obtain the embedding of the whole paragraph⁵ we averaged the word embeddings and passed this vector to the classification head.

To find the best hyperparameters we performed a grid search in each dimension. For the economic dimension, the best hyperparameters consisted of a learning rate of 2e-6, 6 epochs of training, a gamma value of 2, no freezing of the layers, a 768 hidden size, and 10% dropout. For the social dimension, the best hyperparameters were a learning rate of 2e-5, 12 epochs, a gamma of 4, no freezing of the layers, a 768 hidden size, and 10% dropout. Finally, for the foreign dimension the best hyperparameters consisted of a learning rate of 2e-5, 6 epochs, a gamma of 2, no freezing of the layers, a 768 hidden size, and a 10% dropout.

Focal loss. To better address the imbalanced label distribution of this task, we incorporated *focal loss* (Lin et al., 2017), originally proposed for dense object detection. Focal loss can be interpreted as a

	Econ	Social	Foreign	Average
Majority	0.30	0.23	0.25	0.26
BiLSTM	0.44	0.37	0.33	0.38
BERT no-ft	0.46	0.31	0.53	0.44
+pre-training	0.42	0.32	0.46	0.40
BERT ft	0.64	0.50	0.52	0.55
+pre-training	0.56	0.47	0.46	0.49
-focal loss	0.61	0.50	0.50	0.54

Table 6: Macro F1 of the models averaged across 10 runs.

dynamically scaled cross-entropy loss, where the scaling factor is inversely proportional to the confidence on the correct prediction. This dynamic scaling, controlled by hyperparameter γ , leads to a higher focus on the examples that have lower confidences on the correct predictions, which in turn leads to better predictions on the minority classes. Since a γ of 0 essentially turns a focal loss into a cross entropy loss, it has less potential to hurt performance than to improve it. We found the best γ values to be 2 or 4 depending on the dimension.

Task-guided pre-training. We also explored supervised pre-training on two adjacent tasks that can give insights to the relationship between tasks. We used distant supervision that labeled the ideological bias of each article according to that of its news outlet from www.allSides.com (Kulkarni et al., 2018). This procedure allowed us to use the unannotated articles.⁶

6.2 Results

Table 6 shows the macro F1 for each configuration, averaged across 10 runs with different random initializations. The fine-tuned BERT model, with no task-guided pre-training shows the best performance across all 3 ideology dimensions. It is important to note that all the models do better than randomly guessing, and better than predicting the majority class. This shows that the models are capturing some of the complex underlying phenomena in the data. However, the classification tasks still remain challenging for neural models, leaving plenty of room for improvement in future work.

The *BERT ft -focal loss* setting ablates the effect of focal loss against a weighted cross entropy loss, with weights inversely proportional to the distribution of the classes in the dimension. This

⁶We also experimented with pre-training on the dataset from Chen et al. (2020). However, because their dataset starts from 2006 (outside of our time domain), this setting performed poorly.

⁵99% of the paragraphs in the dataset have ≤ 512 tokens.

loss helped get a bump in the macro F1 score of around 0.1 for each dimension compared to an unweighted cross entropy loss. However, the focal loss gave further improvements for 2 of the 3 dimensions. Although task-guided pre-training improved the BERT (no fine-tuning) model for 1 of the 3 dimensions, it led to worse performance than BERT (fine-tuned). The improvement on the no fine-tuning setting indicates that there is a potential correlation to be exploited by the ideology of the news outlet, but such labels are not that informative for multi-dimensional prediction. We hope that this dataset provides a testbed for future work to evaluate more distant supervision data/methods.

7 Related work

In contrast to our multi-dimensional approach that examines the ideology of the issue being discussed instead of the author stance, much of the recent work in computational linguistics has been dedicated to the latter (detection of ideological bias in news media) while collapsing ideology to one dimension (Budak et al., 2016; Kulkarni et al., 2018; Kiesel et al., 2019; Baly et al., 2019, 2020; Chen et al., 2020; Ganguly et al., 2020; Stefanov et al., 2020). The proposed computational models classify the partiality of media sources without quantifying their ideology (Elejalde et al., 2018).

Other researchers interested in the computational analysis of the ideology have employed text data to analyze congressional text data at the legislative level (Sim et al., 2013; Gentzkow et al., 2016) and social media text at the electorate level (Saez-Trumper et al., 2013; Barberá, 2015).

In political science, the relationship between (news) media and polarization is also an active area of research. Prior work has studied media ideological bias in terms of coverage (George and Waldfogel, 2006; Valentino et al., 2009). Prior (2013) argues there is no firm evidence of a direct causal relationship between media and polarization and that this relationship depends on preexisting attitudes and political sophistication. On the other hand, Gentzkow et al. (2016) have established that polarization language snippets move from the legislature in the direction of the media whereas (Baumgartner et al., 1997) have shown that the media has an impact on agenda settings of legislatures.

8 Conclusion

We take the first step in studying multi-dimensional ideology and polarization over time and in news articles relying on the major political science theories and tools of computational linguistics. Our work opens up new opportunities and invites researchers to use this corpus to study the spread of propaganda and misinformation in tandem with ideological shifts and polarization. The presented corpus also provides the opportunity for studying ways that social context determines interpretations in text while distinguishing author stance from content.

This work has several limitations. We only focus on news whereas these dynamics might be different in other forms of communication such as social media posts or online conversations, and the legislature. Further, our corpus is relatively small although carefully annotated by experts. Future work may explore semi-supervised models or active learning techniques for annotating and preparing a larger corpus that may be used in diverse applications.

Acknowledgements

This research is partially supported by NSF grants IIS-1850153, IIS-2107524, and Good Systems,⁷ a UT Austin Grand Challenge to develop responsible AI technologies. We thank Cutter Dalton, Kathryn Slusarczyk and Ilana Torres for their help with complex data annotation. We thank Zachary Elkins, Richard Lau, Beth Leech, and Katherine McCabe for their feedback of this work and its annotation protocol, and Katrin Erk for her comments. Thanks to Pitt Cyber⁸ for supporting this project. We also acknowledge the Texas Advanced Computing Center (TACC)⁹ at UT Austin and The Center for Research Computing at the University of Pittsburgh for providing the computational resources for many of the results within this paper.

References

Alan I Abramowitz and Kyle L Saunders. 1998. Ideological realignment in the us electorate. *The Journal of Politics*, 60(3):634–652.

Theodor Adorno, Else Frenkel-Brenswik, Daniel J

⁷<http://goodsystems.utexas.edu>

⁸<https://www.cyber.pitt.edu>

⁹<https://www.tacc.utexas.edu>

Levinson, and R Nevitt Sanford. 2019. *The authoritarian personality*. Verso Books.

Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational Linguistics*, 34(4):555–596.

Delia Baldassarri and Andrew Gelman. 2008. Partisans without constraint: Political polarization and trends in american public opinion. *American Journal of Sociology*, 114(2):408–446.

Ramy Baly, Giovanni Da San Martino, James Glass, and Preslav Nakov. 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)*, pages 4982–4991.

Ramy Baly, Georgi Karadzhov, Abdelrhman Saleh, James Glass, and Preslav Nakov. 2019. Multi-task ordinal regression for jointly predicting the trustworthiness and the leading political ideology of news media. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2109–2116.

Pablo Barberá. 2015. Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. *Political analysis*, 23(1):76–91.

Frank R Baumgartner, Bryan D Jones, and Beth L Leech. 1997. Media attention and congressional agendas. *Do the media govern*, pages 349–363.

Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623.

Alessandro Bessi, Fabio Petroni, Michela Del Vicario, Fabiana Zollo, Aris Anagnostopoulos, Antonio Scala, Guido Caldarelli, and Walter Quattrociocchi. 2016. Homophily and polarization in the age of misinformation. *The European Physical Journal Special Topics*, 225(10):2047–2059.

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

Ceren Budak, Sharad Goel, and Justin M Rao. 2016. Fair and balanced? Quantifying media bias through crowdsourced content analysis. *Public Opinion Quarterly*, 80(S1):250–271.

Angus Campbell, Philip E Converse, Warren E Miller, and Donald E Stokes. 1980. *The american voter*. University of Chicago Press.

Edward G Carmines, Michael J Ensley, and Michael W Wagner. 2012. Who fits the left-right divide? partisan polarization in the american electorate. *American Behavioral Scientist*, 56(12):1631–1653.

Thomas M Carsey and Geoffrey C Layman. 2006. Changing sides or changing minds? party identification and policy preferences in the american electorate. *American Journal of Political Science*, 50(2):464–477.

Wei-Fan Chen, Khalid Al Khatib, Henning Wachsmuth, and Benno Stein. 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*, pages 149–154.

Dennis Chong and James N Druckman. 2007. Framing theory. *Annual Review of Political Science*, 10:103–126.

Mark Davies. 2012. Expanding horizons in historical linguistics with the 400-million word corpus of historical american english. *Corpora*, 7(2):121–157.

Shrey Desai, Barea Sinno, Alex Rosenfeld, and Junyi Jessy Li. 2019. Adaptive ensembling: Unsupervised domain adaptation for political document analysis. 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 4712–4724.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.

Paul DiMaggio, John Evans, and Bethany Bryson. 1996. Have american’s social attitudes become more polarized? *American journal of Sociology*, 102(3):690–755.

Erick Elejalde, Leo Ferres, and Eelco Herder. 2018. On the nature of real and perceived bias in the mainstream media. *PloS one*, 13(3):e0193765.

Stanley Feldman and Christopher Johnston. 2014. Understanding the determinants of political ideology: Implications of structural complexity. *Political Psychology*, 35(3):337–358.

Morris P Fiorina and Samuel J Abrams. 2008. Political polarization in the american public. *Annu. Rev. Polit. Sci.*, 11:563–588.

Jonathan B Freeman and Rick Dale. 2013. Assessing bimodality to detect the presence of a dual cognitive process. *Behavior research methods*, 45(1):83–97.

Soumen Ganguly, Juhi Kulshrestha, Jisun An, and Hae-woon Kwak. 2020. Empirical evaluation of three common assumptions in building political media bias datasets. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 14, pages 939–943.

Matthew Gentzkow, Jesse M Shapiro, and Matt Taddy. 2016. *Measuring polarization in high-dimensional data: Method and application to congressional speech*. National Bureau of Economic Research.

Lisa M George and Joel Waldfogel. 2006. The new york times and the market for local newspapers. *American Economic Review*, 96(1):435–447.

Jonathan Haidt, Jesse Graham, and Craig Joseph. 2009. Above and below left-right: Ideological narratives and moral foundations. *Psychological Inquiry*, 20(2-3):110–119.

Andrew F Hayes and Klaus Krippendorff. 2007. Answering the call for a standard reliability measure for coding data. *Communication methods and measures*, 1(1):77–89.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.*, 9(8):1735–1780.

Shanto Iyengar, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J Westwood. 2019. The origins and consequences of affective polarization in the united states. *Annual Review of Political Science*, 22:129–146.

Jeffery A Jenkins, Eric Schickler, and Jamie L Carson. 2004. Constituency cleavages and congressional parties: Measuring homogeneity and polarization, 1857–1913. *Social Science History*, 28(4):537–573.

Jacob Jensen, Suresh Naidu, Ethan Kaplan, Laurence Wilse-Samson, David Gergen, Michael Zuckerman, and Arthur Spirling. 2012. Political polarization and the dynamics of political language: Evidence from 130 years of partisan speech [with comments and discussion]. *Brookings Papers on Economic Activity*, pages 1–81.

John T Jost, Christopher M Federico, and Jaime L Napier. 2009. Political ideology: Its structure, functions, and elective affinities. *Annual review of psychology*, 60:307–337.

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.

Vivek Kulkarni, Junting Ye, Steven Skiena, and William Yang Wang. 2018. Multi-view models for political ideology detection of news articles. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3518–3527.

Alban Lauka, Jennifer McCoy, and Rengin B Firat. 2018. Mass partisan polarization: Measuring a relational concept. *American behavioral scientist*, 62(1):107–126.

Yphtach Lelkes. 2016. Mass polarization: Manifestations and measurements. *Public Opinion Quarterly*, 80(S1):392–410.

Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988.

Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.

Roland Pfister, Katharina A Schwarz, Markus Janczyk, Rick Dale, and Jon Freeman. 2013. Good things peak in pairs: a note on the bimodality coefficient. *Frontiers in psychology*, 4:700.

Paul Pierson and Eric Schickler. 2020. Madison’s constitution under stress: A developmental analysis of political polarization. *Annual Review of Political Science*, 23:37–58.

Markus Prior. 2013. Media and political polarization. *Annual Review of Political Science*, 16:101–127.

Martin Riedl and Chris Biemann. 2012. Topictiling: a text segmentation algorithm based on lda. In *Proceedings of ACL 2012 Student Research Workshop*, pages 37–42.

Shruti Rijhwani and Daniel Preotiuc-Pietro. 2020. Temporally-informed analysis of named entity recognition. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7605–7617.

Diego Saez-Trumper, Carlos Castillo, and Mounia Lalmas. 2013. Social media news communities: gatekeeping, coverage, and statement bias. In *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, pages 1679–1684.

Yanchuan Sim, Brice DL Acree, Justin H Gross, and Noah A Smith. 2013. Measuring ideological proportions in political speeches. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 91–101.

Jason Stanley. 2015. *How propaganda works*. Princeton University Press.

Peter Stefanov, Kareem Darwish, Atanas Atanasov, and Preslav Nakov. 2020. Predicting the topical stance and political leaning of media using tweets. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 527–537.

Nicholas A Valentino, Antoine J Banks, Vincent L Hutchings, and Anne K Davis. 2009. Selective exposure in the internet age: The interaction between anxiety and information utility. *Political Psychology*, 30(4):591–613.

Michela Del Vicario, Walter Quattrociocchi, Antonio Scala, and Fabiana Zollo. 2019. Polarization and fake news: Early warning of potential misinformation targets. *ACM Transactions on the Web (TWEB)*, 13(2):1–22.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pier-ric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pages 38–45.

A Data curation

A diachronic corpus is required to measure and analyze polarization over time (DiMaggio et al., 1996). We collect and annotate data across a long period to address the issue of distributional shifts across years (Desai et al., 2019; Rijhwani and Preotiuc-Pietro, 2020; Bender et al., 2021) and help build robust models that can generalize beyond certain periods.

Additionally, the raw data on top of which we annotate needs to satisfy the following constraints: (1) for human annotation to be tractable, the articles should share some level of topical coherence; (2) for the data to be useful for the larger community, the content should also cover a range of common discussions in politics across the aisle; and (3) the articles should come from a consistent set of news outlets, forming a continuous and ideologically balanced corpus.

We start with the diachronic corpus of political news articles of Desai et al. (2019) which covers years 1922-1986, the longest-spanning dataset to our knowledge. This corpus is a subset of news articles from the Corpus of Historical American English (COHA, Davies (2012)). To extract topically coherent articles, we investigate the topics and articles across multiple LDA (Blei et al., 2003) runs varying the number of topics (15, 20, 30, 50), aiming to arrive at a cluster of topics that share common points of discussion and collectively will yield a sizable number of articles each year from the same news outlets.

The LDA models consistently showed one prominent topic—the federal budget—across 5 news outlets with balanced ideology (c.f. Table 8): Chicago Tribune (CT), Wall Street Journal (WSJ), Christian Science Monitor (CSM), the New York Times (NYT), and Time Magazine (TM). Because federal budget stories touch on all aspects of the federal activity, this topic appeals to both liberal and conservative media and thus can provide a good testing ground to showcase our proposed ideological annotation method. In addition to the core federal budget topic (topic 5 of Table 7), we also include other topics such as health and education that are integral parts of ideological beliefs in the United States, and when discussed at the federal government level, are typically related to the federal budget. The top vocabulary of the cluster is shown in Table 7. In an effort to purge articles unrelated to the federal budget, we selected

Topic1: Trade	bank, market, farm, loan, export, agricultur, farmer, dollar, food, debt
Topic2: Business	incom, tax, revenu, profit, corpor, financ, treasuri, pay, sale, bond
Topic3: Education	school, univers, educ, student, colleg, professor, institut, teacher, research, graduat
Topic4: Defense	nuclear, missil, weapon, atom, test, energi, strateg, bomb, space, pentagon
Topic5: Economy	budget, billion, economi, inflat, economic, deficit, unemploy, cut, dollar, rate
Topic6: Health/Race	negro, hospit, medic, health, racial, southern, discrimin, doctor, contra, black
Topic7: Industry	compani, contract, plant, steel, coal, wage, railroad, corpor, manufactur, miner

Table 7: Top words from topics selected in our cluster, from the 50-topic LDA model that yielded the most well-delimited topics.

only those that contain words such as “federal” and “congress”, and excluded those that mention state budget, and letters to editors. (Note that during annotation, we also discard articles that are unrelated to the federal budget.) After this curation, the total number of articles is 5,706 from the 5 outlets.

To account for the sparsity of articles in the first decades and their density in later decades, we narrowed down the articles to the period from 1947 to 1974. We believe this period is fitting because it includes various ideological combinations of the tripartite composition of the American government, Congress and presidency.¹⁰ The total number of articles in the final corpus of political articles on the federal budget from 1947 to 1974 is 1,749.

B Proclaimed ideology of news outlets

	Adfontes	Allsides	MBFC	Average
CSM	-.06	0.00	-.16	-.07
CT	-.04	NA	.34	.15
NYT	-.20	-.5	-.4	-.36
TM	-.10	-.5	-.6	-.4
WSJ	.15	.25	.58	.32

Table 8: Ideological bias of news outlets from common references of media bias. We use the average in our analyses.

¹⁰For example, between 1947-49, Congress was Republican and the President was a Democrat while the story flipped between 1955-57.