

Ideology Prediction from Scarce and Biased Supervision: Learn to Disregard the “What” and Focus on the “How”!

Chen Chen

School of Management
and Economics
CUHK Shenzhen, China
chenchen2020@cuhk.edu.cn

Dylan Walker

Argyros School of
Business and Economics
Chapman University
Dylan@chapman.edu

Venkatesh Saligrama

College of Engineering
Boston University
srv@bu.edu

Abstract

We propose a novel supervised learning approach for political ideology prediction (PIP) that is capable of predicting out-of-distribution inputs. This problem is motivated by the fact that manual data-labeling is expensive, while self-reported labels are often scarce and exhibit significant selection bias. We propose a novel statistical model that decomposes the document embeddings into a linear superposition of two vectors; a latent neutral *context* vector independent of ideology, and a latent *position* vector aligned with ideology. We train an end-to-end model that has intermediate contextual and position vectors as outputs. At deployment time, our model predicts labels for input documents by exclusively leveraging the predicted position vectors. On two benchmark datasets we show that our model is capable of outputting predictions even when trained with as little as 5% biased data, and is significantly more accurate than the state-of-the-art. Through crowd-sourcing we validate the neutrality of contextual vectors, and show that context filtering results in ideological concentration, allowing for prediction on out-of-distribution examples.

1 Introduction

Political Ideology Prediction (PIP) is motivated by a number of applications such as the need to understand government’s policy making (Potrafke, 2018; Hunt, 2009), a legislator’s partisan/non-partisan actions (Zingher, 2014; Nice, 1985; Berry et al., 2007), the general public’s sentiment support for legislation (Berry et al., 2007; Budge et al., 1987; Rudolph, 2009) etc.

Our overall goal is to infer ideology of legislators or the general public from written texts, media posts, and speeches. We can train a PIP model with ground truth labels (e.g., liberal vs conservative) with a standard supervised learning approach. In

some cases, such as for legislators, ground-truth labels can be inferred through party affiliation information, and these have served as good proxies for ideology (Poole and Rosenthal, 2001).

Selection Bias. In other situations, such as, for the domain of social media posts, ground truth is difficult to obtain, and is based on self-reported affiliation of the users (Bakshy et al., 2015). Such self-reported affiliations are generally sparse, and when reported tend to be of extreme polarity (Conover et al., 2011). While manual labeling by annotators (e.g., AMTs) (Preoŕiuc-Pietro et al., 2017) can be leveraged, selection bias is still an issue due to oversampling of posts from extreme vocal posters (Cohen and Ruths, 2013; Mustafaraj et al., 2011; Kohut and Bowman, 2018; Gao et al., 2015).

Justification. A black-box model trained exclusively on scarce and polarized group is likely to perform poorly on the under-observed moderates who are the majority and are reticent to disclose or discuss their politics online (McClain, 2021; Kohut and Bowman, 2018; Gao et al., 2015).

Inferring the majority’s views is important to help us understand support for legislative/executive actions (Craig and Richeson, 2014), or cultural group’s real influence (Bisin and Verdier, 2000) as in the recent Kansas vote where policymakers over-estimated the public support for the abortion ban by overlooking the voice of silent/less vocal majority (Lysen et al., 2022).

Proposed Method. To account for scarcity and selection bias, we propose a novel method that, during training, enforces decomposition of text embedding into context and position, and trains a deep neural network model in a modified variational auto-encoder framework with bi-modal priors.

Contextual Filtering. We propose to decompose document embeddings into a linear (orthogonal) superposition of two *latent* vectors consisting of

a context vector and a position vector. We train a DNN model embedded with this decomposition to enforce the fact that context vectors be devoid of information that contains ideology, and the residual position vectors, obtained by filtering out the context, exclusively bears all ideological information. To ensure that trivial solutions are not found, we require that context vectors across all documents belong to a low-dimensional manifold. Our perspective is that there are a sparse collection of common words (eg. guns, immigration, taxes, etc) whose neutral components serve to contextualize the document, and their corresponding embeddings constitute a low-dimensional space.

De-Noising. Documents exhibit significant variance in scarce regimes, making it difficult to discern ideology. This is due to scarce document collections spread across diverse themes. Position vectors, devoid of context, suppresses noise that accounts for large differences between document embeddings, but carries little relevance for PIP, making learning representation of ideology easier.

Key Experimental Findings

SOTA Prediction. Ideology Prediction by our model BBBG across scarce and biased labels regimes was universally dominant with respect to prior works. Furthermore, compared to other embeddings such as BERT/roBERTa, GloVe predictions were more accurate particularly in scarce/biased regimes.

Neutrality of Latent Contexts. Crowd-sourced experiments showed context vectors are associated with neutral words/phrases across themes.

Ideological Purity. Context Filtering leads to ideological purification and improves prediction.

Knowledge Transfer. BBBG was able to effectively transfer knowledge to out-of-domain scenarios such as (a) generalizing ideology prediction from extremely polarized authors to near-moderate authors; (b) generalizing documents from “seen themes” to documents from “unseen themes.”

2 Related Work

We describe prior works that focus on machine learning for predicting ideology from texts.

Supervised Learning. (Gentzkow and Shapiro, 2010) propose a learning-based approach for ideology prediction on congressional report corpora. Other studies on the same dataset apply modern learning methods under the same setting (Pla and

Hurtado, 2014; Gerrish and Blei, 2011; Iyyer et al., 2014). Research using social media data (tweets, forum posts) aims to map public users onto the liberal-conservative spectrum or simply predict their party labels (Levendusky and Pope, 2010; Baldassarri and Gelman, 2008).

Textual features. The majority of the aforementioned studies utilize off-the-shelf pretrained text representations such as BOW, TF-IDF, LIWC, or Word Embeddings–GloVe, BERT or RoBERTa (Preoțiuc-Pietro et al., 2017; Conover et al., 2011; Makhberian et al., 2020; Liu et al., 2019).

Scarce and Biased Data. The labels used as supervision are obtained from self-reported party affiliations (e.g., Democrats or Republicans) or manual labeling by annotators (Conover et al., 2011; Preoțiuc-Pietro et al., 2017; Cohen and Ruths, 2013), and inherently suffer from label scarcity and selection bias. For example, the majority of the public are reticent to disclose their party affiliation or engage in political discourse online (Bakshy et al., 2015; McClain, 2021). Furthermore, methods proposing to collect labeled texts to study opinions or ideology are prone to over-sampling the “vocal minority” while ignoring or down-sampling the “less vocal majority” (Moe et al., 2011; Kohut and Bowman, 2018; Mustafaraj et al., 2011; Gao et al., 2015). Manual labeling is severely constrained by staffing and cognitive biases of the annotators themselves (Yu et al., 2008; Xiao et al., 2020). Finally, there are large domain differences between content generated by visibly opinionated and moderate users (Cohen and Ruths, 2013). For various reasons, models for learning ideology typically suffer from poor generalizability (Cohen and Ruths, 2013; Yan et al., 2017, 2019).

General Methods. So far there is little research that jointly accounts for both label scarcity and selection bias in training ideology prediction models based on textual data. For Twitter data, social interaction of users (eg. mentions/follows) are leveraged in a graph learning framework to combat the label scarcity issue. However, these methods may not apply to settings where social interaction and connection are absent or unobserved (Xiao et al., 2020). Traditionally, semi-supervised learning (SSL) can deal with insufficient labels (see (Ortigosa-Hernández et al., 2012a; Oliver et al., 2018; Kingma et al., 2014)), whereas Domain Adaptation (DA) (Daumé III, 2009; Glorot et al., 2011) can deal with the distribution discrepancy

between training and test samples, potentially providing a solution to selection bias. We compare these alternative approaches and word-embeddings, and other prior works in our experiments.

3 Statistical Model of Text Generation

Document Decomposition. We first encode documents in a suitable embedding space (GloVe, RoBERTa etc.), and let $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^D$ represent this embedding, which serves as inputs to the learner. We decompose, \mathbf{x} into two major components:

$$\mathbf{x} = \mathbf{c}(\theta) + \mathbf{f}(\theta, \mathbf{z}) + \epsilon \quad (1)$$

where $\mathbf{z} \in \mathbb{R}^M$ is the latent ideology vector for \mathbf{x} , and vectors \mathbf{c} and \mathbf{f} denote the neutral context and filtered position vector components, respectively. ϵ is a random vector that captures idiosyncratic variations among documents. The parameter θ represents the author’s choice of themes such as "guns" or "abortion". However, such decomposition is unidentifiable without further constraints. We impose a low-dimensional structure on context vectors, and during training disambiguate context vectors from ideological content.

Low-Dimensional and Neutral Encoding of Context Vectors. We proceed with the following intuition. For a certain choice of partisan theme θ_i , say that there is a major polarization axis pa_{θ_i} depending on the theme, where people of different ideological groups (such as liberal or conservative, left or right) differentiate when giving speeches¹. In principle we propose to seek orthogonality of context to the polarization axis pa . In practice, since pa is unobserved *ex-ante*, we enforce this constraint by careful initialization and empirically aligning position \mathbf{f} with polarity during training, and we verify the orthogonality *ex-post*.

Initialization. To initialize context vectors, we adopt the following procedure. We first determine themes by generating a set of "neutral" seed words and phrases (hereafter, seeds), either through manual crafting or selecting frequent words with low χ^2 values (see Appendix Sec. A.4), and complement this set by expanding into the neighborhood of seeds, yielding a set of seeds for each theme i , a_{ij} . Themes are initialized as the pooled embedding of seed words $T_{0i} = \sum_j a_{ij}$. Context vectors are initialized as a mixture of themes, $c_0(\theta_0) = T_0\theta_0$. Ul-

¹Mathematically, this can be approximated either via the first principal axis or, in a simpler way, differentiating all average embeddings from the liberal group over the conservative group.

timately, both the theme choice, θ and the themes T are learned through training, starting with the initialization $T = T_0$ (more details and alternate initializations are in the Appendix A.4 and A.5) and low-dimensionality results from the fact that the theme matrix is low-rank. This approach effectively removes variance that accounts for large differences between document embeddings, but carries little relevance for ideology learning and prediction, hence making learning representation of ideology easier.

Multimodal Prior. Because it is commonplace for individuals to explicitly identify (partisanship) with groups of common ideological position, we presume that representation of ideology, \mathbf{z} is drawn from a multi-modal prior. A reasonable choice of modality in the domain of the US politics is two, indicating bipartisanship/polarization. In the context of the congress and political debate forums, this assumption is supported by previous studies (Poole and Rosenthal, 2001; Bonica, 2014) and reflects a general trend in the US society². However, our framework can generalize to multiple modes.

4 Method

Mathematical Background In addition to notation in Sec. 3, we let $\mathbf{y} \in \mathcal{Y}$ denote the output ideological labels taking values in a finite set. Let $p_s(\mathbf{x}, \mathbf{y}), p_t(\mathbf{x}, \mathbf{y})$, denote source and target joint distributions respectively (we drop these subscripts when clear from context). The learner gets access to a labeled set, \mathcal{D}_s of N_s *i.i.d* data points $(\mathbf{x}_i, \mathbf{y}_i) \sim p_s$ and unlabeled set, \mathcal{D}_t , of N_t *i.i.d* target input instances $\mathbf{x}_i \sim p_t$.

Inputs and Outputs of Model. The goal of the learner is to predict labels $\hat{\mathbf{y}}$ for the target given target inputs. Additionally, our trained model also outputs for each input document, the context vector, \mathbf{c} , the theme choices, θ , the position vector, \mathbf{f} , the ideology vector, \mathbf{z} .

Training Loss. Our loss function is a sum of reconstruction loss for inputs $\mathbf{x}_i \in \mathcal{D}_s \cup \mathcal{D}_t$, and cross-entropy loss, $CE(\mathbf{y}_i, \hat{\mathbf{y}}(\mathbf{z}_i; \gamma))$ on $(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{D}_s$; $\hat{\mathbf{y}}$ is the softmax output of a network governed by parameters γ , and taking the encoded latent representation \mathbf{z}_i (see below) as its input.

Reconstruction Loss. Our starting point is to optimize the marginal distribution, namely, $\mathbb{E}_{\mathbf{x} \sim q(\mathbf{x})} \log p(\mathbf{x})$, but as this is intractable, we fol-

²<https://www.pewresearch.org/politics/2014/06/12/political-polarization-in-the-american-public/>

low by relaxing it to the ELBO bound (Kingma and Welling, 2013):

$$L_R(\eta, \phi, \lambda) \triangleq \mathbb{E}_{\mathbf{x} \sim q(\mathbf{x})} [\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\ln p_\eta(\mathbf{x}|\mathbf{z}) + \ln p_\lambda(\mathbf{z}) - \ln q_\phi(\mathbf{z}|\mathbf{x})]] \quad (2)$$

where, $q(\mathbf{x})$ is the empirical distribution on both source and target inputs; $q_\phi(\mathbf{z}|\mathbf{x})$ is the encoder, and $p_\eta(\mathbf{x}|\mathbf{z})$ is the decoder; and $p_\lambda(\mathbf{z})$ is the prior. ϕ, η, λ are their parameters respectively. We now invoke our statistical model Eq. 1 to decompose \mathbf{x} into \mathbf{f} and \mathbf{c} , and by neutrality of context vectors, $p(\mathbf{c}|\theta, \mathbf{z}) = p(\mathbf{c}|\theta)$, and noting that conditioned on θ , \mathbf{c} and \mathbf{f} are independent, we get,

$$L_R(\eta, \phi, \lambda, \theta) \triangleq \mathbb{E}_{\mathbf{x} \sim q(\mathbf{x})} [\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\ln p_\eta(\mathbf{f}|\mathbf{z}, \theta) + \ln p_\eta(\mathbf{c}|\theta) - \ln q_\phi(\mathbf{z}|\mathbf{x}) + \ln p_\lambda(\mathbf{z})]] \quad (3)$$

The first term in Eq. 3 corresponds to the decoder that generates \mathbf{f} ; the second term describes the generation of \mathbf{c} from θ . The last term $p_\lambda(\mathbf{z})$ is the prior distribution of \mathbf{z} . The first three terms minimize the reconstruction error ($\|\mathbf{x} - (\mathbf{f} + \mathbf{c})\|^2$) via mean square loss. Unlike traditional VAE, we enforce multi-modal prior and model it as a K -modal Gaussian following the approach in (Tomczak and Welling, 2018). We approximate our prior by sampling and feeding K pseudo-inputs into the same encoder, namely, $p_\lambda(\mathbf{z}) = \frac{1}{K} \sum_{k=1}^K q_\phi(\mathbf{z}|\mu_k)$. Here K is a hyperparameter that specifies the modality of the prior. For the domain of the US politics we choose $K = 2$, but our method can generalize to settings with different K .

Deep Neural Network (DNN) Training. We train a DNN by optimizing the total loss, $L = L_R(\eta, \phi, \lambda) + \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}_S} CE(\mathbf{y}, \hat{\mathbf{y}}(\mathbf{z}; \gamma))$ end-to-end by backpropagation over all the parameters $(\eta, \phi, \lambda, \gamma)$, while using a bi-modal prior on \mathbf{z} , and call the resulting predictor Bi-Branch Bi-Modal Gaussian Variational Auto-Encoder (BBBG). Our model differs from the traditional VAE in that it uses a multi-modal prior and reconstructs the input using two separately learned components in the generative parts of the model (see Appendix Fig 5, 6).

The Single Branch Ablation. As an ablation, and to understand the importance of Eq. 1, we relax the neutrality constraint by deleting the neutral context learning branch of the model and we call this algorithm SBBG (single-branch variant). In the simplified version, the contexts are no longer estimated and we are reduced to a standard VAE framework with a Gaussian mixture prior, but with

the same K -modal VampPriors. Implementation and training details are described in Sec. 5.5.

Justification for Supervision. Traditional VAE models are unsupervised. But they tend to perform poorly due to over-regularization (Dilokthanakul et al., 2016). In both of our experiments, the unsupervised VAE variant appears to collapse into a prior of uni-modal Gaussian. To solve this issue, we induced the model to converge to a bi-modal prior by providing supervision with a few labeled examples. This means that, during training, we allow *some* ideology labels to be seen by the model, and the prediction loss of ideology are back-propagated to help tune the parameters in the encoders. The number of labels required to effectively train such a system turns out to be 5~8% of the total samples.

5 Experiments

We experiment with two benchmark datasets, *Congressional Speeches* and *Gun Debate Forum* to baseline proposed BBBG, against well-known and prior ML methods. While we provide details of these datasets in the Sec. A.1, we note that the reason for choosing these datasets is driven by the need to ascertain ground-truth extremity of texts and authors (ranging from extreme to neutral). Other datasets are binary, and as such this information is not provided. The datasets we chose allow for sub-sampling of sub-populations with varying levels of ideological extremity. This allows us to validate proposed method under selection bias.

Simulating Label Scarcity and Selection Bias.

We mask ground truth labels in a dataset to simulate label scarcity and selection bias. We define the level of supervision as the percentage of unmasked samples, Sup which determines the extent of masking under either scarcity or scarcity and bias (toward extremity). To simulate scarcity, we masked $(1 - Sup)\%$ of the data randomly. To simulate scarcity and selection bias, we masked $(1 - Sup)\%$ of data from the least extreme authors. We refer to the former procedure as *Unbiased Supervision* and the latter as *Biased Supervision*. Under masking, the prediction loss from masked samples was not used in SGD to update weights of the network. We list details on the masking procedure, and tuning of hyperparameters in Sec. A.2, A.6.

Categorization By Themes. We seek to expose the role of BBBG’s other outputs, namely, the context vector, \mathbf{c} , the filtered position vector, \mathbf{f} , and the ideology vector, \mathbf{z} . To do so, for ablative purposes,

we manually organize (note that BBBG training is agnostic to our categorization) Congressional dataset into 68 themes (see Sec. A.7) consisting of about 25 partisan and 43 non-partisan themes. We perform various experiments to study context neutrality, de-noising through context filtering, and knowledge transfer to unseen themes.

Baseline Methods. Prior methods leverage pre-trained text feature embeddings, and utilize supervised and unsupervised data in various ways. The GloVe embedding is the default embedding for both our main model (BBBG) and other baseline models. Models based on BERT or RoBERTa will be specified through naming. The complete implementation details and comparisons of all models appear in the Sec A.2.

Methods trained only on labeled data. These include GS (Gentzkow and Shapiro, 2010), and standard ML methods such as SVMs, random forests (RFs), XGBoost, and 8-layer Deep Neural Networks (8l-DNN) with similar capacity as BBBG and Gated Recurrent Units. These methods rely on pre-trained text features such as GloVe or RoBERTa (Dey and Salem, 2017; Liu et al., 2019). *Methods leveraging both labeled and unlabeled target data.* These include semi-supervised learning (SSL) methods such as label-spreading with K-nearest neighbor (LS-KNN) (Ortigosa-Hernández et al., 2012b) as well as self-training (ST) methods combined with deep learning (ST-DNN) of similar capacity as BBBG and ensemble learning such as random forests (ST-RF) (Yarowsky, 1995; Tanha et al., 2017; Zhang et al., 2014). ST is based on iterative pseudo-labeling (Zou et al., 2018); and finally Domain Adaptation (DA) methods that are built on RoBERTa embedding (RoBERTa-DA) are applied to handle domain shifts (Glorot et al., 2011).

5.1 Prediction on Congressional Dataset

BBBG outperforms prior works in scarce and biased regimes. The best of baseline models, including deep learning methods, semi-supervised method (LS-KNN or ST-DNN), performed well when outcome labels are abundant and sampling for supervision is unbiased (others such as BERT-DNN, (Devlin et al., 2018) perform poorly (see Sec. B.1)). But with increasing scarcity their performance degrades significantly. As evident in Tab. 1 and Tab. 7 BBBG significantly outperformed all baseline models once the supervision became lower than 20%, and this gap widened with decreasing

supervision for both biased and unbiased supervision. With as little as 5% labels (~10k) with biased sampling for supervision, and 3% (~6.6k) with unbiased supervision, BBBG predictions are highly accurate in predicting party labels of authors.

GloVe vs. BERT and RoBERTa³ First, note that when scarcity/bias is not an issue, BBBG with GloVe vs. other SOTA language models, BERT and RoBERTa (Liu et al., 2019), perform similarly (Tab. 1, Sec. B.1). However, under label scarcity (<8%, which is about 17.6k) and selection bias, RoBERTa embedding performed no better or worse relative to GloVe when combined with either DNN or SBBG framework. This suggests that the more complex model such as RoBERTa or BERT may be more demanding on label abundance and quality and hence more vulnerable to poor supervision. In addition, the Domain Adaption (based on RoBERTa embedding) did not appear to be advantageous than some other baselines, and was significantly below BBBG under almost all conditions. Together, these results show that our BBBG model is considerably more resilient to both label scarcity and selection bias.

5.2 Gun Debate Forum Dataset

BBBG outperforms prior works in scarce, biased and uniformly across all supervision regimes. Under scarcity (<8%, about 4.8k documents) and extremity bias, we compared the performance of BBBG against several alternative models. Apart from those described above we also compared against BERT-Sequential (Devlin et al., 2018; Baly et al., 2020).

BBBG significantly outperforms prior methods on the *Gun Debate Forum* dataset. This is likely due to higher heterogeneity of forum users in their manner of speech compared to Congressional Dataset, which BBBG is able to handle better. We report in Tab. 2 the predicted binary ideology labels (liberal vs conservative) of posters by aggregating predicted outcomes at the author level. Excluding the SBBG ablation, RoBERTa-DA performs the best among the baselines. For SBBG, it only performs well when combining with GloVe embedding instead of RoBERTa. The gap between our model and the best baseline is widened up to 19%

³GloVe has a natural initialization in the form of “single-words” that allows for estimation of contexts and context filtering. We baseline BERT and RoBERTa with SBBG because single-word initialization is difficult as their embeddings are intertwined with surroundings through attention mechanism.

Table 1: Accuracy of party prediction under unbiased/biased supervision for *Congressional Speeches* data, showing competing results between the top three baselines, and the main model BBBG, and the single branch variant SBBG. The best results are in **bold** and the second best are underlined. BBBG outperforms most other models substantially with scarce labels, marked in blue. The percentage shown were averaged over three independent trials.

	Unbiased Supervision								Biased Supervision							
	80%	60%	40%	20%	8%	5%	3%	1%	80%	60%	40%	20%	8%	5%	3%	1%
8I-DNN	93.2%	94.4%	92.6%	88.1%	84.9%	55.0%	53.6%	50.4%	84.5%	84.8%	<u>86.9%</u>	87.9%	65.8%	61.1%	56.7%	50.7%
GS	79.1%	79.9%	81.4%	82.3%	79.0%	77.1%	78.7%	70.4%	70.4%	71.5%	74.5%	72.0%	69.5%	71.1%	68.4%	60.7%
LS-KNN	90.2%	92.6%	92.3%	91.5%	<u>88.8%</u>	85.8%	82.5%	79.7%	<u>84.7%</u>	83.5%	83.0%	81.4%	78.2%	75.7%	71.3%	<u>67.8%</u>
ST-DNN	93.5%	<u>94.3%</u>	94.0%	<u>93.0%</u>	88.0%	77.3%	75.3%	61.6%	81.3%	81.6%	83.2%	80.8%	71.5%	63.8%	56.3%	51.9%
RoBERTa-GRU	92.8%	93.3%	83.3%	78.2%	69.9%	51.6%	51.7%	51.9%	83.6%	84.7%	84.0%	78.1%	53.1%	49.5%	50.7%	50.9%
RoBERTa-DA	90.6%	89.1%	89.4%	79.3%	75.1%	62.9%	61.8%	51.4%	84.0%	81.0%	81.9%	68.5%	69.3%	54.8%	51.4%	49.8%
RoBERTa-SBBG	<u>93.3%</u>	91.1%	<u>94.0%</u>	89.2%	76.5%	61.1%	52.7%	51.3%	86.1%	<u>84.7%</u>	87.5%	<u>85.9%</u>	51.3%	51.2%	48.6%	48.8%
SBBG	92.9%	90.4%	90.6%	89.9%	86.6%	<u>86.7%</u>	<u>84.0%</u>	74.1%	78.3%	81.3%	83.6%	82.0%	<u>79.8%</u>	<u>77.5%</u>	<u>73.5%</u>	65.7%
BBBG	92.7%	92.9%	94.1%	93.2%	91.6%	89.8%	87.2%	81.2%	81.3%	81.3%	85.2%	83.3%	83.6%	84.0%	76.4%	68.6%

Table 2: Accuracy in predicting ideology labels under unbiased/biased supervision for *Gun Debate Forum* data showing competing results between three baselines, the main model BBBG, and the single branch variant SBBG. The best results are in **bold** and the second best are underlined. BBBG outperforms most other models substantially with scarce labels, marked in blue. The percentage shown were averaged over three independent trials.

	Unbiased Supervision								Biased Supervision							
	80%	60%	40%	20%	8%	5%	3%	1%	80%	60%	40%	20%	8%	5%	3%	1%
8I-DNN	72.1%	72.3%	67.4%	61.0%	60.9%	61.2%	61.4%	61.3%	69.9%	69.8%	66.2%	64.8%	61.0%	61.0%	61.2%	<u>61.2%</u>
LS-KNN	66.2%	65.0%	64.3%	64.8%	64.9%	63.9%	62.7%	60.8%	62.4%	60.8%	61.2%	63.4%	61.7%	62.5%	61.3%	61.0%
ST-DNN	70.9%	68.6%	66.8%	64.1%	<60.0%	<60.0%	<60.0%	<60.0%	69.4%	64.8%	64.6%	65.4%	61.9%	61.2%	61.3%	<60.0%
RoBERTa-GRU	76.3%	75.7%	75.9%	71.0%	63.7%	62.2%	61.2%	59.8%	77.6%	73.8%	71.5%	68.9%	67.5%	64.2%	61.1%	58.1%
RoBERTa-DA	72.1%	73.0%	70.3%	66.4%	64.5%	61.5%	62.9%	60.5%	70.3%	73.7%	67.6%	68.9%	68.0%	63.9%	64.1%	60.7%
RoBERTa-SBBG	78.9%	76.7%	73.1%	72.0%	67.8%	62.1%	61.2%	61.1%	74.6%	74.4%	71.0%	69.4%	63.9%	60.7%	60.6%	61.0%
SBBG	96.1%	96.1%	93.6%	91.6%	83.3%	72.4%	65.1%	61.3%	92.2%	88.7%	88.4%	86.1%	73.0%	70.5%	66.6%	60.9%
BBBG	98.8%	97.9%	96.0%	91.6%	85.0%	77.3%	70.7%	61.3%	94.3%	91.0%	90.2%	85.3%	77.3%	74.4%	71.5%	61.2%

difference when the supervision is biased.

5.3 Validation of Context Neutrality

We perform two experimental studies to illustrate the geometry of inferred latent context vectors.

Crowd-Sourced Experiments. We propose to validate, using crowd-sourcing, how well words in the neighborhood of context vectors BBBG outputs aligns with human belief of neutrality. We study:

- 1) Perceived relevance of words to the theme.
- 2) To what extent are these words liberal or conservative (we also include the "do not know" option).

To do so, we calculated the ten closest words to the inferred BBBG context vector of input documents in terms of cosine similarity in the embedded space (Sec. A.8, Tab. 6). We then selected 6 prominent partisan themes (see Tab. 6, A.8) and for each theme randomly sampled 5 out of 10 nearest words mentioned above (hereafter, neighbourhood words). We then chose stereotypical extreme words as references for each theme. These stereotypical extreme words were manually collected by mining five liberal and conservative words/phrases from known partisan news media (see Sec. A.8 and Supplementary). For each item, on Likert Scale, we asked gig workers to rate both relevancy (rescaled to [0,1]) and ideological leaning (rescaled to [-1,1]).

First, we noted that both neighborhood words and manually chosen words were deemed relevant by humans. Words proximal to context

vector scored above 0.843 ± 0.007 in relevance. In comparison, the manually collected conservative and liberal reference words/phrases scored 0.703 ± 0.009 and 0.840 ± 0.007 respectively. On ideological leaning, the neighbourhood words scored 0.056 ± 0.018 while the reference conservative and liberal words/phrases scored 0.323 ± 0.025 and -0.341 ± 0.023 respectively. The neighbourhood words were clearly more neutral (toward 0) than reference words (see Fig. 1). The difference is significant at 0.001 according to a two-sample T-test. Additionally, the Spearman-rho between ranked distance from the context vector and ranked deviation from the neutral point (0) by survey was 0.5, validating that crowd-sourced ratings were aligned with ours (see Sec. A.8).

Orthogonality of Latent Context and Residuals. For each partisan theme (see Sec. A.7), θ , let $\mathcal{D}_\theta/\mathcal{R}_\theta$ denote the data text or speech documents generated by Democrats/Republicans, respectively. We define the *polarization axis* of θ as $\mathbf{pa}_\theta \triangleq \mathbb{E}_{x \sim \mathcal{D}_\theta} \mathbf{f}(x) - \mathbb{E}_{x \sim \mathcal{R}_\theta} \mathbf{f}(x)$, where \mathbf{f} is the output of the filtered position component of BBBG. The angle between context vectors and polarization axis, averaged across different themes was about 84 degrees, which suggests near orthogonality.

5.4 Context Filtering leads to Purity

Here we tested whether upon context filtering, the filtered position vectors, \mathbf{f} , are more concentrated,

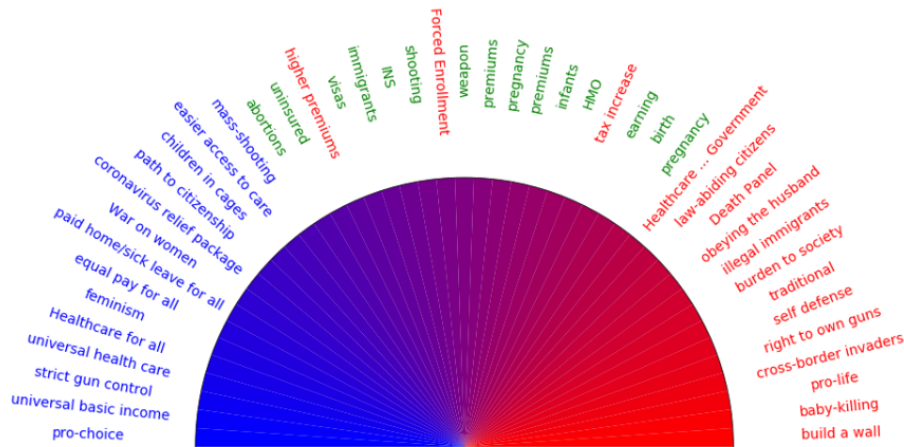


Figure 1: “Sunshine” plot, based on sorting crowd-sourced ratings, depicts how the crowd-workers perceive words in the neighborhood of BBBG context vector (green), relative to other stereotypical words (blue and red) appearing in partisan media. Evidently, these ratings are essentially consistent with BBBG. The few discrepancies appear to be crowd-rating noise.

and exhibit better alignment with ideology labels.

We sampled the speech documents of the most frequent 20 themes given by the top 30 most active speakers (17 Democrats, 13 Republicans) from the Congress (both Senate and House included). For each document, we collected the original input text embedding x as well as corresponding filtered position vector f . We tested two key indicators: (1) % variance explained by the first principal component; (2) the total number of non-zero principal values above given thresholds (i.e. rank of approximated covariance matrix).

Multiple Authors Writings on Diverse Themes.

We explore several variations to highlight denoising effect of BBBG. We report results at a p-value of 0.001:

- 1) *Author writings on one theme:* The % variance explained along the first principal axis increased from 0.22 to 0.39 whereas the ranks decreased from 53 to 14, averaged across all authors and themes.
- 2) *Author writing on multiple themes:* The % variance explained increased from 0.29 to 0.41 whereas the ranks decreased from 101 to 18, averaged across all authors.
- 3) *Multiple authors of same ideology label writing on one theme:* The % variance explained increased from 0.20 to 0.37 for Democrats, and from 0.21 to 0.35 for Republicans, whereas the ranks decreased from 111 to 18 for Democrats, and from 147 to 31 for Republicans, averaged across themes.
- 4) *Multiple authors of same ideology label writing on multiple themes:* the % variance explained increased from 0.26 to 0.40 for Democrats, and from

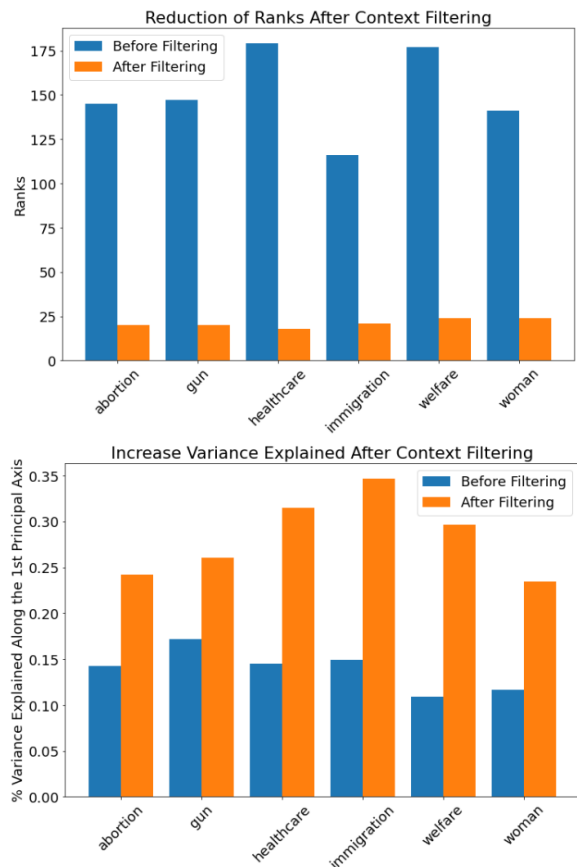


Figure 2: Impact of Context Filtering on Ideology Prediction. (Top) Position Vectors for each theme are concentrated in a low-dimensional space relative to unfiltered encoding. (Bottom) % Variance explained by position vectors is higher relative to unfiltered encoding. These demonstrate significant amplification of ideological signal among a few dimensions.

0.22 to 0.33 for Republicans, whereas the ranks decreased from 151 to 29 for Democrats, and from 147 to 31 for Republicans.

In addition, we found that among each of the six partisan themes, f has significant higher % variance explained at lower ranks (Fig. 2). This demonstrates significant concentration in substantially fewer directions for partisan themes. Together these results consistently suggest context filtering leads to concentration over several "key" directions, and that the inputs used to predict ideology become less "noisy."

The filtered position vectors for different parties are better-separated than original embeddings. Here we sampled the top 20 most frequent themes keeping the other aspects same. In previous analysis we demonstrated that the first principal axis explained at least 35% of total variation from inputs represented as residuals, or 20% for inputs represented as document embeddings. With context filtering the distance between ideological centers⁴ of each party increased 5-fold (from 0.004 to 0.02). Furthermore, when grouped by themes, we observed heterogenous effects, for themes such as "abortion", "guns", "healthcare", and "Iran/Lybia/Syria", the distanced increased significantly (by more than 0.05), while for themes such as "culture", "religion", and "sports", the distances had marginal effect. *Thus themes that are more polarized were better separated along the 1st (i.e. the dominant) Principal Axis.*

Better Ideological Alignment with Labels. To verify whether our model learned a correct representation of ideology, We tested on *Congressional Speeches*, training with 8% biased supervision we observed the following:

- (1) Evidently, ideological members are separated along the 2nd principal component, indicating that z is indeed an ideological representation (Fig 3a).
- (2) Fig 3b shows the mean of the logit output of the ideology supervision module (Fig 6) aggregated by author is highly correlated ($R^2 \sim 0.9$) with *DW-NOMINATE* scores, which is a dataset describing the ideology scores of House and Senate members based on their voting records, and considered a gold-standard (see Sec. A.1). As such *DW-NOMINATE* scores are agnostic to congressional speech, and a high correlation implies BBBG accurately captures ideological differences, compared

⁴Defined as center of normalized vectors projected onto the 1st Principal axis

Table 3: Accuracy in predicting ideology labels under biased supervision for *Gun Debate Forum* data showing competing results between the main model BBBG, and its variants. The best results are in **bold** and the second best are underlined. BBBG outperforms most of its variants substantially with scarce labels, marked in blue. The percentage shown were averaged over three independent trials.

	80%	60%	40%	20%	8%	5%	3%	1%
SBBG_K1	<u>92.4%</u>	<u>90.2%</u>	86.5%	81.5%	73.8%	69.9%	64.7%	60.4%
SBBG_K3	91.0%	89.0%	86.9%	82.3%	73.7%	67.7%	<u>67.9%</u>	59.5%
BBBG_K1	91.9%	90.0%	86.8%	83.8%	<u>77.2%</u>	70.4%	64.1%	58.4%
BBBG_K3	90.5%	88.0%	86.8%	83.3%	76.6%	64.4%	66.1%	59.2%
SBBG	92.2%	88.7%	88.4%	86.1%	73.0%	70.5%	66.6%	60.9%
BBBG	94.3%	91.0%	90.2%	<u>85.3%</u>	77.3%	74.4%	71.5%	61.2%

to baseline models (Fig 3c).

5.5 Results from Ablation Experiments

We tested a few ablations of our main model. The first one is the single branch model (SBBG) where the context learning branch was deleted. Results of this model on both benchmarks were shown in Tab. 1 and 2. As observed in Tab. 1 and Tab. 2, uniformly for all datasets, under either biased or unbiased supervision, BBBG outperforms SSSG. Furthermore, the difference increases at higher scarcity. Since difference in performance in BBBG and SSSG can be attributed to the document decomposition, this implies that the decomposition into neutral context and position vectors results in improved accuracy and generalization to domain shifts.

We also tested impact of K taking values different from 2, such as 1 or 3. Notice that when K=1 the variational part of the model degraded into the ordinary VAE with unimodal Gaussian prior. Tab. 3 showed experiment results of models combining different K values (1 or 3) with SBBG or BBBG (e.g. SBBG_K1 is SBBG with K=1). For comparison purposes, we also included our main model BBBG, and its single branch variants SBBG, both of which has K value equaling to 2. As shown in Tab. 3, when K value deviates from 2, the model performance was worsened. And combining with SBBG would further worsen the performance.

5.6 Knowledge Transfer to Unseen Themes

Better Transfer to Unseen Themes. Here we train samples with supervision labels for documents belonging to eight themes ('IT', 'abortion', 'addictive', 'economy', 'political actions', 'postal service', 'sport', 'traditional energy', 'waste', 'workforce'), which constitute 8% of total documents and contain both partisan and non-partisan themes. We then test on 60 "unseen" themes (see Sec. A.7.

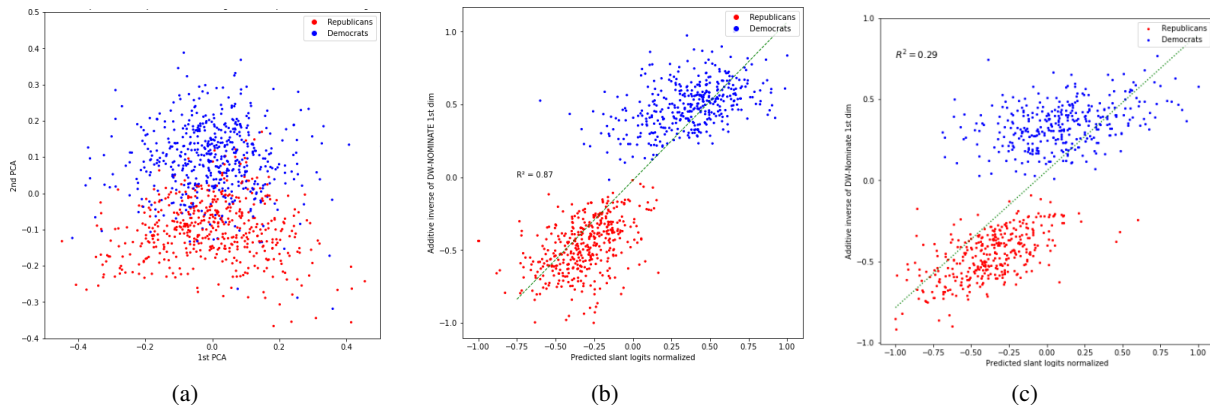


Figure 3: (a) Top 2 PCA components of the ideology vector, \mathbf{z} , suggests that it captures party affiliation. (b&c) Correlation between additive inverse of DW-NOMINATE (y-axis) and slant value (x-axis). (b) BBBG with higher $R^2 \sim 0.9$ correlation. (c) best-performing baseline ($R^2 \sim 0.3$)

Table 4: Experiments on Congressional Reports with supervision labels coming exclusively from samples of eight chosen themes. Here we reported prediction accuracy among document samples of 60 other themes (see Sec. A.7). The best results are in bold and the second best are underlined. BBBG outperforms all other models substantially. The percentage shown were averaged over three independent trials

LS-kNN	RoBERTa -DA	RoBERTa	81-DNN	SBBG	BBBG
63.8%	50.1%	58.2%	78.8%	<u>81.6%</u>	86.1%

As shown in the Tab. 4, BBBG outperformed all other baselines, which demonstrates that context filtered position vectors reveal ideological similarity across different unrelated themes.

BBBG Transfers better from Extreme to Near-Neutral. We evaluate the efficacy of BBBG in generalizing across different biased distributions (See Appendix Fig. 4). Here conservative or extremely conservative posters are about 45% of the posters, and liberal or very liberal are 25%.

This problem is of particular relevance since the majority of the US population is ideologically non-extreme, yet politically inactive (Wojcik and Hughes, 2019), and therefore important to understand (Delli Carpini, 2000). We evaluate our trained models (see Tab. 5) on Gun Debate Forum, and on the sub-population of slightly leaning posters as well as on the held-out set of extreme posters (namely those masked in training). ST-DNN is reported because it is the most representative among prior works, and as observed BBBG’s outperforms ST-DNN. Similar results are observed with other baseline models such as 81-DNN.

Table 5: Comparing prediction accuracy of ideology labels among extreme and slightly leaning groups by models trained on different levels of biased supervision

Supervision Level	Slightly Leaning		Extreme	
	ST-DNN	BBBG	ST-DNN	BBBG
80%	69.3%	83.1%	98.6%	99.2%
60%	63.2%	74.2%	93.4%	99.1%
40%	58.7%	68.2%	88.4%	98.5%
20%	59.4%	65.9%	81.2%	94.5%
8%	58.4%	60.1%	70.4%	85.5%
5%	50.4%	56.5%	70.6%	81.9%

6 Conclusion

We propose a novel deep-learning method to deal with label scarcity and selection bias that typically arise in political ideology prediction problems. Our method learns to decompose text embeddings into neutral contexts and context filtered position vectors, which contain ideological information. In addition to demonstrating improved prediction accuracy on benchmark datasets, we also expose important aspects of our approach through ablative experiments. We validate context neutrality through crowd-sourcing, ideological concentration through context filtering, and knowledge-transfer to documents dealing with novel themes, and to people, whose ideology bears little similarity to training data. Going forward, we can check if our model can extend to more general social platforms such as Twitter, or learn and verify ideological representation on a continuum similar to DW-NOMINATE.

7 Ethical Considerations

7.1 Data Collection and Usage

Congressional Report corpus is publicly available and can be directly downloaded online. Posts from

Debatepolitics.com were collected in a manner that strictly followed the terms of use of the original sources and the privacy rights of the source owner. Authors involved in the data collection process have all finished their human subjects research training.

The Debatepolitics data will be released upon request. The personal information of forum posters will be concealed from the public.

The annotators were recruited from Profic.com and compensated equitably, irrespective of demographic and geographic attributes. Remuneration exceeded the platform's recommended average hourly rate. The study's purpose was disclosed, instructions were clear, and the task posed no harm to participating annotators.

7.2 Benefit and Potential Misuse of BBBG

Intended Use. The goal of this project is to provide a means to overcome the label bias and scarcity problem that has not been fully addressed in the ideology prediction literature. It also provides a useful representation of ideology that can be further explored for other learning purposes. It is particularly useful to predict and evaluate the stance of the non-extreme group who tends to politically inactive (cf. sec 4.2.3).

Recent Kansas abortion vote⁵ has demonstrated the importance of predicting leanings of the silent majority. Devoid of such tools, lawmakers are more likely to incorrectly extrapolate views of the vocal minority to the entire population. Furthermore, poor extrapolation emanating from the vocal minorities views can have a significant impact on political disengagement⁶.

However, like any machine learning model, there is a risk of over-generalization of its true capabilities. The output of our model needs to be assessed and evaluated with full consideration of the characteristics of the input source. The potential domain difference of authors of texts might be significant, and any conclusion drawn from studying our group of authors cannot be immediately generalized to other groups.

Risk of Misuse and Potential Harm. Our model should not cause harm unless its users interpret the prediction results in an unintended way.

⁵<https://apnews.com/article/kansas-abortion-vote-recount-e874f56806a9d63b473b24580ad7ea0c>

⁶https://academic.oup.com/pa/article/68/suppl_1/241/1403570

It is meant to provide insights on the ideology distribution of a group or a population, instead of judgment of an individual. Its output is not without error, albeit more accurate than most models under realistic situations. And for slightly leaning and moderate people, it is possible our model may generate incorrect outputs relative to the ground truth. Though, our model mitigates this relative to the prior SOTA. The potential harm of our model could be magnified if it is used in making decisions on vulnerable populations.

The predictions and insights generated by our model should not be treated as facts or golden rules. We also suggest that results from any political related studies should be interpreted with skepticism and encourage the users of our model to perform careful evaluation in their corresponding application domain, check more sources or consult political scientists for expert opinions.

Acknowledgements This research was supported by the Army Research Office Grant W911NF2110246, AFRL Grant FA8650-22-C-1039, the National Science Foundation grants CCF-2007350 and CCF-1955981, and the Hariri Data Science Faculty Fellowship Grant.

References

- Eytan Bakshy, Solomon Messing, and Lada A Adamic. 2015. Exposure to ideologically diverse news and opinion on facebook. *Science*, 348(6239):1130–1132.
- 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. *arXiv preprint arXiv:2010.05338*.
- William D Berry, Evan J Ringquist, Richard C Fording, and Russell L Hanson. 2007. The measurement and stability of state citizen ideology. *State Politics & Policy Quarterly*, 7(2):111–132.
- Alberto Bisin and Thierry Verdier. 2000. A model of cultural transmission, voting and political ideology. *European Journal of Political Economy*, 16(1):5–29.
- Adam Bonica. 2014. Mapping the ideological marketplace. *American Journal of Political Science*, 58(2):367–386.
- Ian Budge, Hearl Derek, David Robertson, Derek Hearl, et al. 1987. *Ideology, strategy and party change: spatial analyses of post-war election programmes in 19 democracies*. Cambridge University Press.
- Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794.
- Raviv Cohen and Derek Ruths. 2013. Classifying political orientation on twitter: It’s not easy! In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 7.
- Michael D Conover, Bruno Gonçalves, Jacob Ratkiewicz, Alessandro Flammini, and Filippo Menczer. 2011. Predicting the political alignment of twitter users. In *2011 IEEE third international conference on privacy, security, risk and trust and 2011 IEEE third international conference on social computing*, pages 192–199. IEEE.
- Maureen A Craig and Jennifer A Richeson. 2014. On the precipice of a “majority-minority” america: Perceived status threat from the racial demographic shift affects white americans’ political ideology. *Psychological science*, 25(6):1189–1197.
- Hal Daumé III. 2009. Frustratingly easy domain adaptation. *arXiv preprint arXiv:0907.1815*.
- Michael X Delli Carpini. 2000. Gen. com: Youth, civic engagement, and the new information environment. *Political communication*, 17(4):341–349.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Rahul Dey and Fathi M Salem. 2017. Gate-variants of gated recurrent unit (gru) neural networks. In *2017 IEEE 60th international midwest symposium on circuits and systems (MWSCAS)*, pages 1597–1600. IEEE.
- Nat Dilokthanakul, Pedro AM Mediano, Marta Garnelo, Matthew CH Lee, Hugh Salimbeni, Kai Arulkumar, and Murray Shanahan. 2016. Deep unsupervised clustering with gaussian mixture variational autoencoders. *arXiv preprint arXiv:1611.02648*.
- Guodong Gao, Brad N Greenwood, Ritu Agarwal, and Jeffrey S McCullough. 2015. Vocal minority and silent majority. *MIS quarterly*, 39(3):565–590.
- Matthew Gentzkow and Jesse M Shapiro. 2010. What drives media slant? evidence from us daily newspapers. *Econometrica*, 78(1):35–71.
- Sean M Gerrish and David M Blei. 2011. Predicting legislative roll calls from text. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*.
- Xavier Glorot, Antoine Bordes, and Yoshua Bengio. 2011. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *ICML*.
- Geoffrey Hinton, Nitish Srivastava, and Kevin Swersky. 2012. Neural networks for machine learning lecture 6a overview of mini-batch gradient descent. *Cited on*, 14(8):2.
- Michael H Hunt. 2009. *Ideology and US foreign policy*. Yale University Press.
- Mohit Iyyer, Peter Enns, Jordan Boyd-Graber, and Philip Resnik. 2014. Political ideology detection using recursive neural networks. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1113–1122.
- John E Jackson and John W Kingdon. 1992. Ideology, interest group scores, and legislative votes. *American Journal of Political Science*, pages 805–823.
- Diederik P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. 2014. Semi-supervised learning with deep generative models. In *Advances in neural information processing systems*, pages 3581–3589.
- Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Andrew Kohut and Carol Bowman. 2018. The vocal minority in us politics. In *Radio—The Forgotten Medium*, pages 45–58. Routledge.

- Matthew S Levendusky and Jeremy C Pope. 2010. Measuring aggregate-level ideological heterogeneity. *Legislative Studies Quarterly*, 35(2):259–282.
- Steven D Levitt. 1996. How do senators vote? disentangling the role of voter preferences, party affiliation, and senator ideology. *The American Economic Review*, pages 425–441.
- Shuhua Liu and Patrick Jansson. 2017. City event identification from instagram data using word embedding and topic model visualization. *Working Papers*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Dylen Lysen, Laura Ziegler, and Blaise Mesa. 2022. Voters in kansas decide to keep abortion legal in the state, rejecting an amendment. *NPR News*.
- Colleen McClain. 2021. [70% of u.s. social media users never or rarely post or share about political, social issues](#). *Pew Research Center*.
- Wendy W Moe, David A Schweidel, and Michael Trusov. 2011. What influences customers’ online comments. *MIT Sloan Management Review*, 53(1):14.
- Negar Mokherian, Andrés Abeliuk, Patrick Cummings, and Kristina Lerman. 2020. Moral framing and ideological bias of news. In *International Conference on Social Informatics*, pages 206–219. Springer.
- Eni Mustafaraj, Samantha Finn, Carolyn Whitlock, and Panagiotis T Metaxas. 2011. Vocal minority versus silent majority: Discovering the opinions of the long tail. In *2011 IEEE Third International Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third International Conference on Social Computing*, pages 103–110. IEEE.
- David C Nice. 1985. State party ideology and policy making. *Policy Studies Journal*, 13(4):780.
- Avital Oliver, Augustus Odena, Colin Raffel, Ekin D Cubuk, and Ian J Goodfellow. 2018. Realistic evaluation of deep semi-supervised learning algorithms. *arXiv preprint arXiv:1804.09170*.
- Jonathan Ortigosa-Hernández, Juan Diego Rodríguez, Leandro Alzate, Manuel Lucania, Inaki Inza, and Jose A Lozano. 2012a. Approaching sentiment analysis by using semi-supervised learning of multi-dimensional classifiers. *Neurocomputing*, 92:98–115.
- Jonathan Ortigosa-Hernández, Juan Diego Rodríguez, Leandro Alzate, Manuel Lucania, Inaki Inza, and Jose A Lozano. 2012b. Approaching sentiment analysis by using semi-supervised learning of multi-dimensional classifiers. *Neurocomputing*, 92:98–115.
- Ferran Pla and Lluís-F Hurtado. 2014. Political tendency identification in twitter using sentiment analysis techniques. In *Proceedings of COLING 2014, the 25th international conference on computational linguistics: Technical Papers*, pages 183–192.
- Keith T Poole and Howard Rosenthal. 2001. Denominate after 10 years: A comparative update to congress: A political-economic history of roll-call voting. *Legislative Studies Quarterly*, pages 5–29.
- Niklas Potrafke. 2018. Government ideology and economic policy-making in the united states—a survey. *Public Choice*, 174(1):145–207.
- Daniel Preotiuc-Pietro, Ye Liu, Daniel Hopkins, and Lyle Ungar. 2017. Beyond binary labels: political ideology prediction of twitter users. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 729–740.
- Juan José Rodríguez, Ludmila I Kuncheva, and Carlos J Alonso. 2006. Rotation forest: A new classifier ensemble method. *IEEE transactions on pattern analysis and machine intelligence*, 28(10):1619–1630.
- Thomas J Rudolph. 2009. Political trust, ideology, and public support for tax cuts. *Public Opinion Quarterly*, 73(1):144–158.
- Suzanna Sia, Ayush Dalmia, and Sabrina J Mielke. 2020. Tired of topic models? clusters of pretrained word embeddings make for fast and good topics too! *arXiv preprint arXiv:2004.14914*.
- Jafar Tanha, Maarten van Someren, and Hamideh Afzarmanesh. 2017. Semi-supervised self-training for decision tree classifiers. *International Journal of Machine Learning and Cybernetics*, 8(1):355–370.
- Jakub Tomczak and Max Welling. 2018. Vae with a vampprior. In *International Conference on Artificial Intelligence and Statistics*, pages 1214–1223. PMLR.
- Stefan Wojcik and Adam Hughes. 2019. Sizing up twitter users. *PEW research center*, 24.
- Zhiping Xiao, Weiping Song, Haoyan Xu, Zhicheng Ren, and Yizhou Sun. 2020. Timme: Twitter ideology-detection via multi-task multi-relational embedding. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2258–2268.
- Hao Yan, Sanmay Das, Allen Lavoie, Sirui Li, and Betsy Sinclair. 2019. The congressional classification challenge: Domain specificity and partisan intensity. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 71–89.
- Hao Yan, Allen Lavoie, and Sanmay Das. 2017. The perils of classifying political orientation from text. In *LINKDEM@ IJCAI*.

- David Yarowsky. 1995. Unsupervised word sense disambiguation rivaling supervised methods. In *33rd annual meeting of the association for computational linguistics*, pages 189–196.
- Bei Yu, Stefan Kaufmann, and Daniel Diermeier. 2008. Classifying party affiliation from political speech. *Journal of Information Technology & Politics*, 5(1):33–48.
- Pengyuan Zhang, Yulan Liu, and Thomas Hain. 2014. Semi-supervised dnn training in meeting recognition. In *2014 IEEE Spoken Language Technology Workshop (SLT)*, pages 141–146. IEEE.
- Joshua N Zingher. 2014. The ideological and electoral determinants of laws targeting undocumented migrants in the us states. *State Politics & Policy Quarterly*, 14(1):90–117.
- Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. 2018. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European conference on computer vision (ECCV)*, pages 289–305.

A Appendix

A.1 Dataset Details.

Congressional Speeches is a corpus of transcriptions of (220k) speeches given by house or senate congresspeople during 2009-2020, spanning both the Obama and Trump presidencies. For each speech, the speaker ID, year, party affiliation of the speaker, and the title of the speech was provided. The ideological label of each speech document is given as the party affiliation of its speaker. Although such a proxy seems imperfect, it has been shown that party affiliations are largely consistent with congresspeople’s ideology (Levitt, 1996). Extremity of documents/speakers can be ascertained from DW-Nominate scores.

Gun Debate Forum is a corpus of 60K posts on the online forum (debatepolitics.com) debating the issues of gun violence and gun control. Each user posting in this forum may choose to affiliate with one of the following ideological groups: slightly liberal, slightly conservative, liberal, conservative, extreme liberal, extreme conservative, moderate, centrist, libertarian, anarchist, and progressive. In this study, we limited our attention to the liberal-conservative spectrum, including posts from users who identified as: slightly liberal, slightly conservative, liberal, conservative, extreme liberal, extreme conservative, and moderate (Fig. 4). This places all posts and their authors on a 7-point scale (Preoțiuc-Pietro et al., 2017).

DW-NOMINATE is a dataset describing the ideology scores of House and Senate members based on their voting records, obtained from (voteview.com). Scores (primary dimension) range continuously from liberal (-1) to conservative (+1) and explain the majority (>80%) of voting differences. Details of how scores are calculated from voting data are provided in (Poole and Rosenthal, 2001). DW-NOMINATE scores are widely considered as a benchmark metric (Jackson and Kingdon, 1992).

A.2 Baseline Methods and Experiment Details

81-DNN A fully-connected (FC) network that has 7 intermediate layer and an output layer of size 1. For GloVe embedding, the input document embedding is obtained by averaging word embeddings with the size of 300, while for BERT, the input document is represented using CLS (or pooled) embedding with the size of 1024. The intermediate layers are of size (in order) 800, 800, 800, 400, 250, 800, and 800 (to reduce verbosity, the shape of such FC network

will be written as (800, 800, 800, 400, 250, 800, 800, 1), same for other FCs hereinafter). We used ReLU as activation function except for the output layer where we used Sigmoid function. We used 0.001 learning rate for training, l2-regularization at 0.01 at the last two layers, and RMSProp for optimization (Hinton et al., 2012). The 8l-DNN will henceforth be used as building blocks for some other baseline models. Unless specified, it will use the same structure and parameters as above.

The GS model originated from the benchmark method measuring political slants from texts, developed by Gentzkow and Shapiro in 2020 (Gentzkow and Shapiro, 2010). It is based on the Naïve Bayes assumption. We repeated what described in (Gentzkow and Shapiro, 2010) by picking out most polarized phrases, and then regressing party labels over word frequency, and then use the sum of coefficients of those polarized phrases weighted by phrase frequency to obtain the party slant of each speech in the test data.

GRU This model takes texts as sequences of input embedding (Dey and Salem, 2017), and output a vector of length 300. This output vector was further fed into a Dropout layer ($p=0.5$), then into a FC network of a shape (800, 800, 400, 1). All layers use ReLU activation except for the last layer which uses Sigmoid. To train this GRU, we used ADAM optimizer with starting learning rate of 0.01

The BERT/roBERTa model Both the BERT and roBERTa model output two types of representation for sentences - pooled/CLS embeddings or sequences of embeddings (Devlin et al., 2018; Baly et al., 2020; Liu et al., 2019). We tried both in our experiments and only reported results from sequential representations from roBERTa as it produced the best performance. The sequence of embeddings were fed as input into a network of the same structure described in the GRU. For the roBERTa-SBBG model, the sequence of embeddings were first fed as input into a single GRU layer to generate embeddings of size 300, which were subsequently treated as input samples for a framework described in SBBG model⁷. Unlike GloVe, we allowed the embeddings to be fine-tuned during training.

⁷We only tested roBERTa with the SBBG ablation instead of BBBG full model due to the fact that BERT/roBERTa embeddings are intrinsically non-linear which are created from weighted sum of all contextual information, and hence conflicted from the linear decomposition model of documents which is the backbone of BBBG. However, as shown in our experiments, switching to roBERTa from GloVe offered little benefits in improving the learning performance.

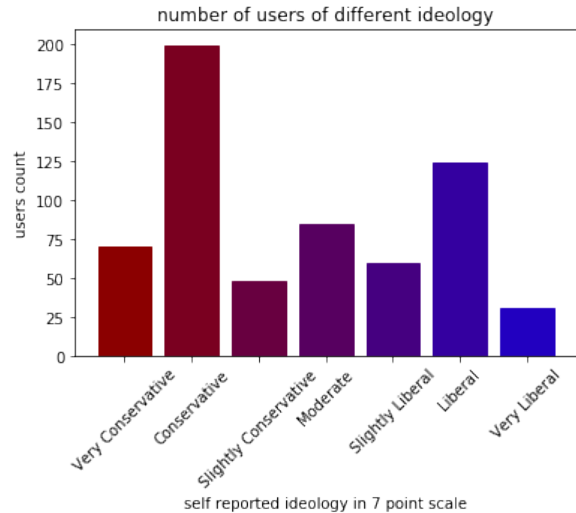


Figure 4: Number of posters per ideology group. Observe that conservative or extremely conservative posters are about 45% of the posters, and liberal or very liberal are 25%. The situation is reversed among slightly leaning posters. They are of slightly liberal than slightly conservative.

Domain Adaptation The Domain Adaptation module was based on sequence embeddings produced by roBERTa. During the training stage, the labeled texts were assigned to the source domain, and equal number of randomly sampled unlabeled texts were assigned to the target domain. During the inference stage, predictions from the label classifier were reported on the test data.

Label Spreading and Self Training are both prototypical semi-supervised learning (SSL) models to deal with situations in which the training data are not sufficiently labeled (Yarowsky, 1995; Ortigosa-Hernández et al., 2012b). Although supervision insufficiency is a predominant obstacle in ideology learning, SSL is rarely used in this field of study and was included to see why it was not widely adopted to address this issue. In our experiments, we used kNN ($k=7$) kernel for Label Spreading, as well as 8l-DNN, Random Forest or XGBoost as kernels for Self Training (Rodriguez et al., 2006; Chen and Guestrin, 2016).

A.3 BBBG Training

According to eq3, the generative part of BBBG model can be visualized as in fig5, while the whole model is illustrated in fig6. It contains four major components, listed as followed.

\mathbf{x} to \mathbf{z} . The encoder of the VAE framework, where the mean and standard deviation of \mathbf{z} ($\text{dim}=50$) is learned from \mathbf{x} through two fully-




Figure 5: A diagram of generative part of BBBG.




Figure 6: Proposed bi-branch BBBG network with VampPriors. Each arrow here represented multiple layers (>3), excluding the reconstruction layer. The green square highlights the ideology supervision module.

connected (FC) networks. z is sampled through reparametrization trick. The shape of FC to learn the mean of z is (800, 800, 800, 400, 50). The FC to learn the variance is of the same shape, and shares the weight with the mean network at the first two intermediate layers. All layers use ReLU activation function except for the last layer of the variance FC, which is a step function that maps 1 for all inputs between -6 to 3 but 0 elsewhere.

z to y. The ideology prediction element is composed of a FC network of a shape (800, 800, 1). It uses ReLU as the activation function, except for the last layer where it uses Sigmoid.

x to $\tilde{\theta}$ to c. This is the neutral context branch. The theme assignment vector $\tilde{\theta}$ is learned from x through a FC network of the shape (800, 1600, 400, 68), where 68 is the total number of themes (including "others"). The context vector c is obtained by multiplying the theme matrix T and $\tilde{\theta}$ (explained below).

z and $\tilde{\theta}$ to f. The decoder of the VAE framework that decodes f . z and $\tilde{\theta}$ are first concatenated. Then the resulting vector will serve as the input for a FC network of a shape (800, 800, 300) to decode f . Next \tilde{X} will be reconstructed by sum of f and c .

The generation of priors of z (VampPriors) from pseudo-inputs follows as in (Tomczak and Welling, 2018).

Hyperparameters For experiment on Congressional Report, we tuned hyperparameters for our main model (the single branch version inherited

the same set of hyperparameters) on 10% of total samples as the validation sets. We used 0.001 learning rate with the RMSProp Optimizer. Other parameters included the mean (0.02) and standard deviation (0.3) of Gaussian distribution where VampPriors were sampled from (the value in the bracket indicating the final tuned value)(Tomczak and Welling, 2018), the annealing parameter of KL-divergence of the z prior (15)(Dilokthanakul et al., 2016), the number of training epochs (10 for experiment 1, 15 for experiments 2). We used L2-regularization at 0.01 in the ideology prediction branch. Experiment on Debatepolitics corpus inherited the same hyperparameters as above.

SBBG Same as the BBBG but without x to $\tilde{\theta}$ to c , also x is directly reconstructed from z , with a FC network of a shape (800, 800, 300).

A.4 Initialization of Theme matrix, T and theme assignment θ

While theme assignment can be obtained by unsupervised models like the latent dirichlet allocation (LDA), we may encounter two issues. First, LDA models might uncover non-neutral associations. Second, a model purely based on word distribution might fail to uncover infrequent themes. In congressional reports, themes are dominated by fiscal spending, healthcare, economy, and politics ($>60\%$). Themes that were significant yet infrequent, such as guns ($<1\%$) and abortion ($<2\%$), could easily be missed by LDA models. To address these points, we adopted a popular embedding similarity approach(Liu and Jansson, 2017; Sia et al., 2020). First, we manually chose 67 major themes manually by digesting a significant collection of the congressional report (5K) and summaries of 5k bills; for each theme, we constructed a set of neutral words, by asking experts to come up with 5 to 10 neutral seed words (such as "firearms" and "guns" for the gun theme), and expanding into 100 words by the proximity in the word embedding space; next we manually checked each seed and expanded seed word to remove the irrelevant, rare or "non-neutral" words (i.e., words that are clearly partisan, such as "Obamacare" for the healthcare theme, "sinful" for the LGBTQ theme, and "baby-killing" for the abortion theme); then the embedding of a theme is calculated by averaging embeddings of all chosen words; lastly all theme vectors obtained above were stacked in an order to form the initial theme matrix T_0 , of which each row corresponded

to a theme vector of a certain theme, and would be subsequently used to initialize the theme vector matrix T (an alternative way to decide sets of neutral words via statistical inference can be found in the next session). If we know theme assignment of a document x , then the context vector is $c = T\theta$. For Debatepolitics, the theme assignment θ is known. For Congressional Reports, the theme assignment was initialized by taking cosine similarity between document embeddings and context vectors obtained above (Liu and Jansson, 2017; Sia et al., 2020), with a slight modification. Instead of obtaining a hard assignment, we generated a initial soft assignment θ_0 by taking a Softmax transformation of embedding of one document with each of the context vectors. We denoted the theme as "other" if the maximum cosine similarity falls below a certain threshold (around 10% quantile of all maxima).

We then used these initial theme assignment in the training process, as described in the main paper, thereby allowing both θ and T to be updated and fine-tuned with x . Shown in Fig. 6, $\hat{\theta}$ were learned from x through a feedforward network, and c was obtained by $c = \hat{T}\hat{\theta}$. Here \hat{T} was initialized with values of T_0 , and both $\hat{\theta}$ and \hat{T} are trainable variables. In the loss function, two additional terms were included to minimize the mean-squared difference between $\hat{\theta}$ (or \hat{T}) and the initial values of θ_0 (or T_0). This constrains the search for the local optima of θ and T around the neighborhood of their original values.

When training was finished, we performed post hoc verification of the updated neutral context vector matrix T , by 1) checking whether it was orthogonal to the polarization axis corresponding to each theme (see main paper); 2) manually verified the neighborhood of the context vector for each theme in the original embedding space. For each theme, we collected the top 20 closest words to the context vectors in the embedding space.

A.5 An Alternative Way of Obtaining Neutral Words for the Neutral Context Vectors

This section will offer a statistical definition of theme-wise neutral words which can be used as an alternative way to initialize the neutral context vectors. Consider a choice of theme, such as "guns", let \mathcal{W}_{gun}^D be the set of words or phrases commonly used by Democrats to discuss about guns, and \mathcal{W}_{gun}^R be the set of words or phrases

commonly used by Republicans. For each word $w \in \mathcal{W}_{gun} \triangleq \mathcal{W}_{gun}^D \cup \mathcal{W}_{gun}^R$, we calculated its document frequency fr_{doc} , as well as the discrepancy of document frequency among different party fraction $\Delta fr \triangleq |fr_{doc}^D - fr_{doc}^R|$ (or alternatively, χ^2 of each words (Gentzkow and Shapiro, 2010)). Words w whose $fr_{doc}(w) \leq \alpha$ and $\Delta fr(w) \leq \beta$ were eliminated from \mathcal{W}_{gun}^D and \mathcal{W}_{gun}^R . Finally the neutral set of words for the gun theme is defined as $\mathcal{W}_{gun}^N \triangleq \mathcal{W}_{gun}^D \cap \mathcal{W}_{gun}^R$. This neutral set contains words such as "gun", "guns", "ammunition", etc.

By this definition, words that are clearly ideologically driven (such as "libtards", "baby-killing", "death (panel)" etc.) are removed. In addition, the neutral set rules out words related to a certain theme that were preferred by a certain ideological group, such as "illegal" vs "undocumented" of the immigration theme. This allows the residual to capture as much information as possible to make the correct inference of the ideology. This approach is inspired by Gentzkow and Shapiro's paper, where they adopted a similar procedure, but different from us chose the most ideological words with highest χ^2 values.

A.6 Additional Details on Masking

The masking procedure was performed either by random sampling or by selecting the top X% (X ranging from 80 to 1) of the most extreme Democrats and Republicans as unmasked. Here the extremity was determined using congress members' DW-NOMINATE score, which is considered as a benchmark metric for political ideology (Jackson and Kingdon, 1992). We refer to the former type of supervision as unbiased supervision (as the sample will contain both extreme and non-extreme members) and the latter type as biased supervision.

For the second experiment, the representation of texts as word embeddings and the masking procedures were the same as in experiment 1. However, due to the fact that the ground truth ideology is represented not continuously but on a 7-point scale, and the fact the distributions of users and posts from each group were uneven, we used a weighted sampling scheme **without** replacement to simulate the observed outcome scarcity of less extreme populations, as follows.

In the weighted sampling scheme, the posts generated by extreme groups (i.e. extremely liberal or conservative) were 10 times more likely to be sampled into the unmasked portion compared to the

regular group (i.e. liberal or conservative), which were subsequently 20 times more likely to be sampled into the unmasked portion compared to the slightly leaning group (slightly liberal or conservative). In this way, when the unmasked samples were scarce compared to the total population (less than 8%), they will predominantly consist of the posts generated by the extreme posters. And even when the coverage of unmasked samples reaches 80% of the population, they are still very unlikely to include samples from the slightly leaning users. This scheme mimics the real world scenario where more extreme individuals are more politically vocal and more willing to disclose their own ideology, whereas the "silent majority", who are mostly moderate or slightly leaning, are relatively nonvocal politically, and their political views and ideology labels are largely unobserved.

A.7 All and Partisan Themes for Congressional Reports

The names of all themes are "IP", "IT", "abortion", "academic", "agriculture", "business", "children", "China", "commerce", "crime", "culture", "homeland security", "detainee", "disadvantaged", "disaster", "race and minorities", "disease", "addictive drugs", "economy", "education", "environment", "family", "federal operation", "finance", "fiscal themes", "food", "gun", "health theme", "healthcare", "high tech", "housing", "immigration", "industry", "cyber-security", "infrastructure", "international", "Iran Syria Libya", "Iraq", "Israel", "jury", "lgbtq", "media", "military complex", "natives", "nuclear", "police", "political actions", "postal", "R&D", "religion", "renewable energy", "reserves", "Russia", "safety", "sport", "tax", "terrorism", "trade", "traditional energy", "transportation", "veteran", "vietnam", "vote", "waste", "welfare", "woman", "workforce", and "other".

The polarized themes selected to verify context vectors' neutrality (excluding uncommon themes that occurred less than 500 times) are "abortion", "agriculture", "detainee", "disadvantaged", "disaster", "race and minorities", "disease", "addictive drugs", "economy", "environment", "fiscal", "gun", "health themes", "healthcare", "immigration", "international", "Iran Syria Libya", "Iraq", "Israel", "military complex", "renewable energy", "Russia", "traditional energy", "welfare", "workforce".

A.8 Additional Details on Crowd-Sourced Experiments

To validate how well learned context vectors from BBBG aligns with human belief of neutrality, we recruited 476 participants on the Prolific platform, all of whom are citizens of the United States. Among those, the majority of them living in California (67), New York (40), and Texas (34). 50% of the surveyees are female and 47% are male. About 55% of the participants have bachelor's or higher degrees, 34% have high school diplomas, and 10% have community college degrees. As for occupations, more than a quarter of participants report they are currently unemployed. Among the participants who are employed, most of them are working in management and professional jobs (129), followed by service (89) and sales (53). According to self-reported ideology on a seven-point scale, the majority (62%) support Democratic viewpoints, while participants who support Republican views and those who are moderates/centrists account for 19% each. Each participant received equal compensation at an hourly rate of 9.22 USD, surpassing the average rate recommended by the platform.

We extracted context vector of six popular partisan themes (Abortion, Gun, Healthcare, Immigration, Social Welfare, and Women's Right), and calculated cosine distance ($1 - \text{cosine similarity}$) between context vectors and projected embedding of each of top 10K words in the lexicon. The 5 of 10 closest words for each theme was selected as *context neighborhood* words. As controls, we manually mined five stereotypical conservative/liberal words/phrases from well-established partisan news media (later we describe sources of each word/phrase). Those will be referred as *reference* words. Hence, 15 words/phrases will be surveyed for each theme (75 in total).

During the survey, each surveyee was asked to complete several demographic questions, including his/her own ideology leaning (liberal, conservative, center/neither), as well as scoring on 30 randomly sampled target or reference words/phrases. The instructions were as follows: "The following pages will provide you with some topics and keywords. You will have to determine whether the keywords are relevant to the topic and to what extent they are associated with liberals or conservatives." More specifically, we asked them to rate 1) to what extent (from 1 to 7 on the Likert scale) they believed those words were related to each theme, and




Figure 7: Average ideological leaning scores for the Abortion theme according to crowd workers. The target, liberal, conservative words/phrases were colored in green, blue, and red respectively. y axis is ideological leaning score. The color of each point also corresponds to the value of the score.

2) to what extent they believed those words were leaning to the liberal end or the conservative end (1 being very liberal and 7 being very conservative; an "do not know/unsure" option was also included to increase accuracy of the rating). To avoid subjective biases (for example, liberal surveyees tend to rate liberal words closer to neutral), we randomly down-sampled the results given by liberal surveyees for each words/phrases (since there are approximately twice as many liberal crowd workers as conservatives or centrists on Prolific). Eventually each word/phrase was rated by approximately the same number (around 65 on average) of liberals, conservatives, and centrists. During analysis, we rescaled the relevancy scores to $[0, 1]$ and ideological leaning scores to $[-1, 1]$.

We calculated the ideological leaning score for each words, and aggregated them to the theme level (see Tab. 8. Fig7 showed the ideological leaning scores of surveyed words/phrases (3 out of 5 were chosen for display clarity) for the Abortion theme.

A.9 Sense-making \mathbf{z} : Does it Capture Ideology?

If \mathbf{z} is ideological in nature, then there should exist some principle axes along which members of common ideological groups (party labels) are differentiated. Fig 3a shows the top two PCA components of the (mean aggregated by author) ideology \mathbf{z} for all authors, with colors corresponding to author party affiliation (Republicans red; Democrats blue). With few exceptions, members are well separated along the 2nd principal component, indicating that \mathbf{z} is indeed an ideological representation.

If our model produces an accurate representation and estimation of the true ideology of authors, then the mean of output aggregated by author should be highly correlated with traditional ideology scores obtained from data independent from *Congressional Speeches*, such as the *DW-NOMINATE* scores based on voting behavior. Fig 3 shows the relationship between our model's predicted slant score, which is the output of the machinery predicting ideological labels from \mathbf{z} , and that of the primary DW-NOMINATE score for authors. The two scores are highly correlated ($R^2 \sim 0.9$), indicating that our model accurately estimated the ideology of document authors with minimum supervision. Having established the validity of our model, we now turn to evaluations in the face of label scarcity and extremity bias.

B Additional Results

B.1 Additional Table for Experiment 1

Table 7 provides additional competing results between BBBG and several other baselines mentioned in the main paper, performed over the congressional report corpus. This is to provide evidence that methods tabulate in main paper outperform methods reported here.

B.2 BBBG can discriminate among slightly leaning groups even when this information is absent in training data.

BBBG and prior works report meaningful performance with sufficient supervision. Therefore, a key insight of this experiment is that in heterogeneous group settings, when there is even less data for slightly leaning groups because of biased supervision, a model that performs well must be particularly capable of extracting ideological knowledge.

We compared the different schemes using a novel rank-deviation (RD) metric (see supplementary) that compares the ranking of an author in the ideological spectrum against the median rank. As evident, the average or median RDs of the slightly leaning subgroup are smaller than RDs of the extreme subgroup (see Fig. 8) BBBG achieves better separation between extreme and slightly leaning subgroups and lower variance, compared with the ST-DNN.

Table 6: Selected reference and target words/phrases for survey experiments on Prolific.

	ref_Conservative	ref_Liberal	Context_Neighborhood
Abortion	Right to Life, baby-killing, immoral, partial-birth abortion, pro-life	abortion right, pro-choice, My Body My Choice, reproductive right, women's right	abortion, abortions, birth, pregnancy, pregnant
Gun	2nd Amendment, gun right, law-abiding citizens, self-defense right to own guns	gun ban, gun safety, mass-shooting, school massacre, strict gun control	firearms, gun, guns, shooting, weapon
Healthcare	Death Panel, Forced Enrollment, Healthcare in the hand of Government Bureaucrats, higher premiums, tax increase	Affordable care, Healthcare for all, easier access to preventive and primary health care, expanding Medicaid universal health care	Health Maintenance Organization, healthcare coverage, Medicaid, premiums, uninsured
Immigration	alien criminals, Build A Wall, traffickers, illegal immigrants cross-border invaders	children cages, legalization of immigrants, moratorium on deportations, path to citizenship, undocumented immigrants coronavirus relief package.	Immigration and Naturalization Service (INS), aliens, immigrants, refugees, visas
Welfare	benefit for the lazy, burden to the society, opportunity society, too big government, welfare queen	paid home/sick leave for all workers, social safety net, universal child care, Universal Basic Income (UBI)	earning, incomes, premiums, salary, wages
Women's Right	family values, loyalty to marriage, obeying the husband, traditional, virtual of women	War on women, equal pay for equal work, feminism, girls, me too	babies, infants, pregnancy, teenagers, women

Table 7: Accuracy of party prediction under unbiased or biased supervision for *Congressional Speeches* data, showing competing results between other baselines and the main model. The best results are in **bold** and the second best are underlined. BBBG outperforms most other models substantially with scarce labels, marked in blue. The percentage shown were averaged over three independent trials.

	Unbiased Supervision								Biased Supervision							
	80%	60%	40%	20%	8%	5%	3%	1%	80%	60%	40%	20%	8%	5%	3%	1%
ST-RF	70.0%	71.5%	71.1%	69.7%	69.7%	69.8%	72.8%	61.1%	74.3%	<u>72.9%</u>	<u>68.9%</u>	56.9%	61.3%	61.3%	61.3%	61.3%
BERT	67.4%	67.3%	66.4%	60.4%	53.8%	50.4%	50.0%	51.2%	64.7%	65.4%	66.1%	57.7%	64.8%	61.0%	61.0%	61.2%
RF	65.1%	64.8%	64.6%	64.6%	64.4%	63.6%	63.0%	<u>61.2%</u>	61.7%	57.5%	54.2%	53.0%	63.4%	61.7%	62.5%	61.3%
GRU	77.7%	<u>77.3%</u>	<u>75.7%</u>	<u>72.4%</u>	69.2%	67.5%	65.7%	60.3%	61.6%	57.3%	54.1%	52.9%	86.1%	<u>73.0%</u>	<u>70.5%</u>	<u>66.6%</u>
BBBG	92.7%	92.9%	94.0%	93.2%	91.6%	89.8%	87.2%	81.2%	81.3%	81.3%	85.2%	83.3%	85.3%	77.3%	74.4%	71.5%

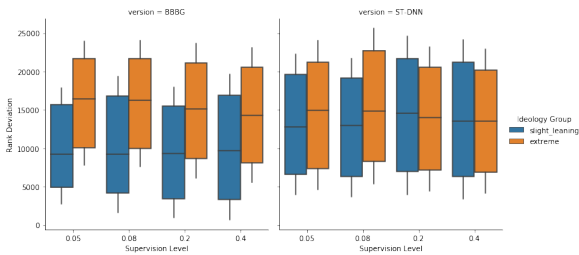


Figure 8: Comparing the distributions of Rank Deviation of various subgroup, according to BBBG (left) and ST-DNN (rank), trained on different levels of supervision (percentage along X-axis).

B.3 Additional Results on Crowd-Sourced Experiments

Detailed comparison of crowd-rated ideological leaning score by each theme can be found in Tab. 8.

Table 8: Ideological leaning of words according to crowd-sourced experiments, breaking down by themes. Each gig worker was asked to rate randomly sampled words that belong to our targeted neighborhood words or one of the reference group (Liberal words or Conservative words). Scores are re-scaled to the range of -1 to 1, where -1/1 corresponds to extreme liberal/extreme conservative leaning.

Themes	References		Targets
	Conservative	Liberal	Neighborhood Words
Abortion	0.524*** (0.082)	-0.507*** (0.062)	0.039 (0.052)
Economy	0.166 (0.074)	-0.341*** (0.073)	0.093 (0.05)
Finance	0.047 (0.073)	-0.125*** (0.077)	0.058 (0.046)
Gun	0.498*** (0.067)	-0.338*** (0.079)	0.157 (0.067)
Healthcare	0.097* (0.081)	-0.413*** (0.06)	-0.0 (0.052)
Immigration	0.514*** (0.077)	-0.223** (0.083)	-0.041 (0.069)
Iraq	0.317*** (0.063)	-0.126*** (0.071)	0.108 (0.055)
Renewable	0.246*** (0.074)	-0.439*** (0.053)	-0.038 (0.05)
Welfare	0.377*** (0.094)	-0.493*** (0.055)	0.124 (0.047)
Woman	0.418*** (0.07)	-0.364*** (0.074)	0.08 (0.049)

The significance of crowd rated ideological leaning difference between our target words and either one of the reference group (Liberal or Conservative) was calculated via Two sample T-Test. The significance level was indicated behind reference groups.

* p<0.05; ** p<0.01; *** p<0.001

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Left blank.
- A2. Did you discuss any potential risks of your work?
Left blank.
- A3. Do the abstract and introduction summarize the paper’s main claims?
Left blank.
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
No response.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
No response.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
No response.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
No response.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
No response.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
No response.

C Did you run computational experiments?

5

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Appendix A.1, A.2, A.3

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

No response.

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

sec 5, appendix B

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

For details implementing RoBerta and Sklearn, see appendix A.2

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Sec 5 and A.8

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

A.8 for instructions. The experiments involves evaluating neutrality of a few words, there is no risk.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

sec7 and A.8

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

sec 7 and A.8

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

sec 7 and A.8

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

A.8