

Party Matters: Enhancing Legislative Embeddings with Author Attributes for Vote Prediction

Anastassia Kornilova Daniel Argyle Vlad Eidelman

FiscalNote

{anastassia.kornilova,daniel,vlad}@fiscalnote.com

Abstract

Predicting how Congressional legislators will vote is important for understanding their past and future behavior. However, previous work on roll-call prediction has been limited to single session settings, thus did not consider generalization across sessions. In this paper, we show that metadata is crucial for modeling voting outcomes in new contexts, as changes between sessions lead to changes in the underlying data generation process. We show how augmenting bill text with the sponsors' ideologies in a neural network model can achieve an average of a 4% boost in accuracy over the previous state-of-the-art.

1 Introduction

Quantitative analysis of the voting behavior of legislators has long been a problem of interest in political science, and recently in NLP as well (Gerrish and Blei, 2011; Kraft et al., 2016). One of the most popular techniques in political science for modeling legislator behavior is the application of spatial, or ideal point, models built from voting records (Poole and Rosenthal, 1985; Clinton et al., 2004), that are often used to represent uni-dimensional or multi-dimensional ideological stances. While *roll call votes* (i.e Congressional voting records) provide explanatory power about a legislators position with respect to previously voted-on bills, these models are limited to in-sample analysis, and are thus incapable of predicting votes on new bills.

To address this limitation, recent work has introduced methods that take advantage the text of the bill, along with the voting records, to model Congressional voting behavior (Gerrish and Blei,

2011; Nguyen et al., 2015; Kraft et al., 2016). This work is related to a long line of studies on using political text to model behavior, ranging over political books, Supreme Court decisions, speeches and Twitter (Mosteller and Wallace, 1963; Thomas et al., 2006; Yu et al., 2008; Sim et al., 2016; Iyyer et al., 2014a; Sim et al., 2013; Preoțiu-Pietro et al., 2017).

In addition to enabling prediction, associating text with ideology allows for a further degree of interpretability. However, all previous work incorporating text into roll call prediction have limited their evaluation to in-session training and testing.¹

As legislators typically serve for multiple sessions, and similar bills are proposed across sessions, we want to be able to leverage this data across sessions to inform our model. However, the generalizability of previous methods to a cross-session setting is unknown.

In this work, we explore the problem of roll call prediction across sessions. We show that previous methods are unable to generalize across sessions, thus suggesting that current text representations are not sufficient for modeling voting outcomes in new contexts. We hypothesize that each session has a different underlying data generation process, wherein the ideological position of the observed bills varies depending on the controlling party. This is supported by the observation that about 75% of bills up for a vote in a given session have a sponsor in the party in power.

As noted in Linder et al. (2018), the policy area, or topic, of the bill, and the ideological position, are two separate dimensions underlying the text. Since legislators tend to sponsor bills that are ideologically aligned with them, a model trained on a single session will mostly be exposed to bills with a specific ideology on each topic. Thus, a

¹A session is a 2-year period of legislative business.

single session model may get the ideology information as an implicit prior without needing to explicitly capture it. This challenge was not obvious in previous studies that were limited to a single session. Across sessions, however, the ideological prior on a given topic changes, resulting in variations in voting patterns that are not captured by current text modeling methodologies alone.

In applications where the text may contain an insufficient signal, researchers may turn to additional metadata features. This technique has previously been used in various contexts, such as incorporating sponsor and committee features for predicting bill committee survival (Yano et al., 2012), and enhancing tweet recommendations with location data (Xing and Paul, 2017).

We propose a neural architecture that directly models the ideological variation across sessions using metadata about the bill sponsors, and show that this can strongly improve performance with little overhead to complexity and training time.

2 Model

Spatial voting models assume that a legislator has a numeric *ideal point* which represents their ideology. They make voting decisions on bills, which also have a numeric representation. While the details of the implementation vary,² spatial voting models share the idea that the closer a bill’s representation is to a legislator’s ideal point the more likely the legislator is to vote *yes*.

Following this framework, we model the core vote prediction problem as follows: Given a legislator, L , and a bill, B , predict their vote y , with possible outcomes: *yes* or *no*.

Using these inputs, let v_L be an embedding representing the legislator, and v_B be the bill embedding. First, v_B is projected into the legislator embedding space:

$$v_{BL} = \mathbf{W}_B v_B + \mathbf{b}_B \quad (1)$$

where \mathbf{W}_B and \mathbf{b}_B are a weight matrix and a bias vector, respectively. Then, we measure the alignment between the two vectors. Previous work used a dot-product for this step, instead, we express the comparison as follows:

$$\mathbf{W}_v (v_{BL} \odot v_L) + \mathbf{b}_v \quad (2)$$

²For example, Poole and Rosenthal represent bills as cut-points that divide legislators into yes and no groups (Poole and Rosenthal, 1985) and later work based on item response theory conceptualizes bills as “discrimination” vectors that are multiplied by an ideal point vector.

where \odot represents element-wise multiplication, and \mathbf{W}_v is a weight vector of the same dimensions as v_L . Finally, we apply a sigmoid activation function to get the vote prediction:

$$p(y = \text{yes} | B, L) = \sigma(\mathbf{W}_v (v_{BL} \odot v_L) + \mathbf{b}_v) \quad (3)$$

Using this architecture, we develop several novel bill representations. First, we consider different text-only representations, then we show how to incorporate metadata.

2.1 Text Model

Previous work incorporating text has primarily been based on topic models (Gerrish and Blei, 2011; Lauderdale and Clark, 2014; Nguyen et al., 2015) and embeddings (Kraft et al., 2016). As the embedding framework achieved superior performance, we adopt a similar architecture. While Kraft et al. (2016) represented the text using a mean word embedding (MWE) representation, we replace it with a Convolutional Neural Network (CNN) representation (Kim, 2014), which has achieved superior performance on recent text classification tasks (Dauphin et al., 2016; Wen et al., 2016; Yang et al., 2016). Our CNN uses 4-grams and 400 filter maps.

2.2 Sponsor Metadata

We posit that a legislator’s voting behavior is influenced both by the topic and the ideology of a bill. A legislator may be more liberal on one issue and more conservative on another. Thus, we need to capture both aspects. While previous work has shown that text alone contains ideological information (Iyyer et al., 2014b), the metadata of the bill may be a stronger source, especially for ideology. This approach has had success in the related problem of bill committee survival,³ where signals about the sponsors, committee and chamber were used in conjunction with text models (Yano et al., 2012).

We use this idea to improve our bill representations. One particularly strong signal is the author of the bill, because of their ideological motives. For simplicity, we represent the bill’s authorship as the percentage of Republican and Democrat sponsors (p_r and p_d). We propose that the Republican and Democratic sponsors influence the text

³Congressional bills, first, are voted on in a committee, before moving to the floor.

of the bill in different ways. To obtain the overall ideological position of the bill, we combine the versions of the bill influenced by each party. The final bill can thus be represented as follows:

$$v_B = ((\mathbf{a}_r p_r) \cdot T_r) + ((\mathbf{a}_d p_d) \cdot T_d) \quad (4)$$

where T_r and T_d are the Republican and Democratic copies of the text representation (e.g MWE or CNN); p_r and p_d are the scalars representing the percentage of sponsors from each party (e.g 0.7 and 0.3); and \mathbf{a}_r and \mathbf{a}_d are vectors representing how the percentages should influence each dimension of the text embedding.

The larger p_r or p_d is, the stronger the influence of that party on the bill.

We test two text representations for T_r and T_d : one using MWEs and one using CNNs. The underlying word embeddings are initialized with 50d GloVe vectors (Pennington et al., 2014) and are non-static during training.

The rest of the model weights are initialized randomly with the *glorot uniform* distribution (Glorot and Bengio, 2010). The length of v_L is set to 25. All models are trained using binary cross-entropy loss and optimized with the *AdaMax* algorithm (Kingma and Ba, 2014). The models are trained for 50 epochs, using mini-batches of size 50.

3 Dataset

Our dataset was collected from GovTrack,⁴ and consists of nonunanimous roll call votes and texts of resolutions and bills introduced in the 106th to 111th Congressional sessions.⁵ We also collect the bill summaries written by the Congressional Research Service⁶ (a non-partisan organization), that provide shorter descriptions of the key actions in each bill. All text is preprocessed by lowercasing and removing stop-words.

As bills are often much longer than the typical document encountered in other NLP tasks, with an average of 2683 words per bill, and some bills having hundreds of pages, with correspondingly

⁴<https://theunitedstates.io/>

⁵We exclude bills with unanimous votes because these are typically associated with routine matters (for example, the naming a post office or an official commendation) that do not contain ideological motivation. We consider bills where less than 1% of legislators voted ‘no’ to be unanimous; about 42% of bills fall into this category.

⁶<https://www.congress.gov/help/legislative-glossary/>

Session	Total Bills	Total Votes	% Yes Votes
2005-2012	1718	685,091	68.4%
2013-2014	360	136,807	66.4%
2015-2016	382	153,605	61.8%

Table 1: Count of Bills and Votes

Session	House Majority	Senate Majority
2005-2006	R	R
2007-2008	D	D
2009-2010	D	D
2011-2012	R	D
2013-2014	R	D
2015-2016	R	R

Table 2: Party in power by session

lengthy summaries, this poses a problem for our compositional neural architecture. To address this, we limit the length of each full-text and summary to N words, where N is empirically set to the 80th percentile of the collection. For summaries $N=400$, and for full-text $N=2000$.

4 Experiments

As described earlier, the experimental framework in previous work treated each session individually. To evaluate the ability of our model to generalize across sessions, we perform several sets of experiments. In the first set, in-session, we perform 5 fold cross-validation over the 2005-2012 sessions. In the second, out-of-session, we train on multiple sessions, 2005-2012, and evaluate on sessions not included during training, the 2013-2014 and 2015-2016 sessions. During testing, we only include legislators present in the training data.

The overall statistics for our dataset are presented in Tables 1 and 2.

5 Results

To understand how sponsor parties and text interact in the input, and how our predictive power changes when testing on in-session bills and out-of-session bills. We test the following models:

- MWE: mean word embedding text model as described in Kraft et al. (2016) using summaries;
- MWE+FT: MWE model using full bill text;

- CNN: text model from Section 2.1 over summaries;
- MWE+Meta: MWE representation combined with metadata as described in Section 2.2;
- CNN+Meta: like MWE+Meta but using a CNN instead of averaging;
- MWE+Meta+FT: As above using full bill text;
- Meta-Only: A variation on MWE+Meta that uses the same, random “dummy” text for all the bills, only changing the metadata (p_r and p_d).

Each model is first evaluated in-session, where both train and test bills come from the same set of sessions, and thus same distribution, and then out-of-session, where training bills are from one set of sessions and the model is evaluated on a different set. All results are presented in Table 3.

5.1 In-session Results

We evaluate our models with accuracy on 5-fold cross-validation. All three models combining text with metadata perform significantly better than the others, showing that the text and meta information have complimentary predictive power, and that our models’ sponsor-augmented text representation is able to capture the ideological preference. The CNN+Meta achieves the highest accuracy of 86.21, followed by MWE+Meta at 85.96, showing that the CNN learns a somewhat better text representation than MWE. Compare this to the baseline MWE model without meta information, which achieves an accuracy of 81.10, only slightly better than the Meta-Only model at 80.27. Contrary to our hypothesis, MWE achieves higher accuracy than Meta-Only. However, it remains unclear whether this signal is related to ideology or other contextual information. The performance on the out-of-session setting will determine whether this signal is akin to ideology.

5.2 Out-of-session Results

In this setting, on both test sessions, text with meta information achieves the best performance as well. On the 2013-2014 session, the CNN+Meta model does the best at 83.59. Unlike the in-session setting, Meta-only does better than the text-only

	in-session	out-of-session	
	2005-2012	2013-2014	2015-2016
Guess Yes	68.31	65.92	61.07
MWE	81.10	77.57	69.80
MWE + FT	81.46	68.33	57.94
CNN	83.24	77.49	69.63
Meta-Only	80.87	82.28	67.10
MWE + Meta	85.96	82.73	71.90
MWE+Meta+FT	85.14	82.43	69.86
CNN + Meta	86.21	83.59	70.99

Table 3: Accuracy Results

models (MWE, CNN). This supports the theory that within the sessions we are able to capture contextual ideology from the text, but once we move to a new session the text models no longer contain an accurate representation of the Congressional ideology.

While in other experiments we are able to achieve at least a 17% improvement over the Guess Yes baseline, on 2015-2016, the best model, MWE+Meta, is only able to achieve a 10% gain. During this session divisions arose within the Republican party in the House of Representatives that disrupted the typical voting dynamics.⁷ Unlike 2013-2014, the Meta-Only model does worse than the text ones; however, the gap between them is much smaller.

5.3 Overall Analysis

These experiments provide several interesting insights. First, because using both text and metadata (MWE+Meta or CNN+Meta) results in the strongest model in every case, we confirm that legislators vote based on both the topic and the ideology of the bill.

Second, the text-only models do significantly worse on the out-of-session tests than the in-session ones. This confirms our theory that session-specific contextual information is implicitly captured by the previous single-session models, but that context is not accurate in new sessions. If we were capturing ideology from the text, then the text only model should have performed well out-of-session.

⁷A conservative bloc of the Republican Party (the Freedom Caucus) began to assert influence over party leadership, eventually resulting in the ouster of John Boehner as Speaker (Lizza, 2017).

Third, to further examine whether a neural model was the best technique for modeling text with metadata, we trained a SVM model over the bag-of-words representation of the summary, indicator variables for the legislators and the percent of bill sponsors in each party (e.g p_d). This model did not perform as well as either MWE or Meta-Only, showing that the embedding approach is better at representing this combination of features.

Finally, the models that embed the full text (+FT) generally perform worse than embedding the summaries. While this confirms that the summary contains sufficient information about the topics and the actions in the bill, we did not fully explore the bill text.

6 Future Work

While Congress introduces close to 20,000 bills every session, very few of them receive a vote, limiting the dataset. We would like to explore various bootstrapping techniques that would allow us to expand the dataset size with artificial votes.

Furthermore, while our text representations are sufficient for modeling shorter text, i.e. summaries, we would like to test more sophisticated representations in the future, in particular, those designed to handle longer texts.

7 Conclusion

In this paper, we developed a neural network architecture to predict legislators votes that augments bill text with sponsor metadata. We introduced a new evaluation setting for this task: out-of-session performance; which allows us to examine the generalizability of our proposed model, and was not considered in past studies. Finally, we showed that the introduction of metadata to bias the text representations outperforms the existing text-based methods in all experimental settings.

References

Joshua Clinton, Simon Jackman, and Douglas Rivers. 2004. The statistical analysis of roll call data. *American Political Science Review*, 98(2).

Yann N. Dauphin, Angela Fan, Michael Auli, and David Grangier. 2016. [Language modeling with gated convolutional networks](#). *CoRR*, abs/1612.08083.

Sean Gerrish and David M Blei. 2011. Predicting legislative roll calls from text. In *Proceedings of ICML*.

Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of AISTATS*.

Mohit Iyyer, Peter Enns, Jordan Boyd-Graber, and Philip Resnik. 2014a. Political ideology detection using recursive neural networks. In *Proceedings of ACL*, volume 1, pages 1113–1122.

Mohit Iyyer, Peter Enns, Jordan Boyd-Graber, and Philip Resnik. 2014b. [Political ideology detection using recursive neural networks](#). In *Association for Computational Linguistics*.

Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv:1408.5882*.

Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *CoRR*, abs/1412.6980.

Peter Kraft, Hirsh Jain, and Alexander M Rush. 2016. An embedding model for predicting roll-call votes. In *Proceedings of EMNLP*.

Benjamin E. Lauderdale and Tom S. Clark. 2014. Scaling politically meaningful dimensions using texts and votes. *American Journal of Political Science*, 58(3):754–771.

Fridolin Linder, Bruce A. Desmarais, Matthew Burgess, and Eugenia Giraudy. 2018. Text as policy: Measuring policy similarity through bill text reuse. *SSRN/2812607*.

Ryan Lizza. 2017. [The war inside the republican party](#).

Frederick Mosteller and David L Wallace. 1963. Inference in an authorship problem: A comparative study of discrimination methods applied to the authorship of the disputed federalist papers. *Journal of the American Statistical Association*, 58(302):275–309.

Viet-An Nguyen, Jordan L. Boyd-Graber, Philip Resnik, and Kristina Miler. 2015. [Tea party in the house: A hierarchical ideal point topic model and its application to republican legislators in the 112th congress](#). In *Proceedings of ACL*.

Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of EMNLP*, pages 1532–1543.

Keith T Poole and Howard Rosenthal. 1985. A spatial model for legislative roll call analysis. *American Journal of Political Science*, pages 357–384.

Daniel Preotiu-Pietro, Ye Liu, Daniel Hopkins, and Lyle Ungar. 2017. Beyond binary labels: political ideology prediction of twitter users. In *Proceedings of ACL*.

Yanchuan Sim, Brice DL Acree, Justin H Gross, and Noah A Smith. 2013. Measuring ideological proportions in political speeches. In *Proceedings of EMNLP*.

- Yanchuan Sim, Bryan R. Routledge, and Noah A. Smith. 2016. Friends with motives: Using text to infer influence on scotus. In *Proceedings of EMNLP*.
- Matt Thomas, Bo Pang, and Lillian Lee. 2006. Get out the vote: Determining support or opposition from congressional floor-debate transcripts. *CoRR*, abs/cs/0607062.
- Ying Wen, Weinan Zhang, Rui Luo, and Jun Wang. 2016. Learning text representation using recurrent convolutional neural network with highway layers. *CoRR*, abs/1606.06905.
- Linzi Xing and Michael J Paul. 2017. Incorporating metadata into content-based user embeddings. In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pages 45–49.
- Xiao Yang, Craig MacDonald, and Iadh Ounis. 2016. Using word embeddings in twitter election classification. *CoRR*, abs/1606.07006.
- Tae Yano, Noah A Smith, and John D Wilkerson. 2012. Textual predictors of bill survival in congressional committees. In *Proceedings of NAACL*. Association for Computational Linguistics.
- Bei Yu, Stefan Kaufmann, and Daniel Diermeier. 2008. Classifying party affiliation from political speech. *Journal of Information Technology & Politics*, 5(1):33–48.