Learning Alignments from Legislative Discourse

Daniel Kauffman
California Polytechnic State University
San Luis Obispo, California, USA
dkauffma@calpoly.edu

Foaad Khosmood
California Polytechnic State University
San Luis Obispo, California, USA
foaad@calpoly.edu

Toshihiro Kuboi
Institute for Advanced Technology and Public Policy
San Luis Obispo, California, USA
tkuboi@calpoly.edu

Alex Dekhtyar
California Polytechnic State University
San Luis Obispo, California, USA
dekhtyar@calpoly.edu

ABSTRACT

In this work, we seek to quantify the extent to which a legislator's spoken language indicates their degree of alignment toward an organization that has a taken a documented position on some legislation. To perform this study, we use a corpus of bill discussion transcripts provided by Digital Democracy[1]. We then apply proven learning methods in the field of natural language processing to predict alignment scores between each member of the California state legislature and a select set of state-recognized organizations. Our methods surpass established baselines, achieving up to 78% accuracy when predicting these same scores using discourse features.

CCS CONCEPTS

• **Applied computing** → **Computing in government**; • **Computing methodologies** → *Machine learning approaches*; *Model development and analysis*;

KEYWORDS

natural language processing, government transparency, political influence

ACM Reference Format:
Daniel Kauffman, Foaad Khosmood, Toshihiro Kuboi, and Alex Dekhtyar. 2018. Learning Alignments from Legislative Discourse. In *dg.o '18: dg.o 2018: Proceedings of the 19th Annual International Conference on Digital Government Research, May 30-June 1, 2018, Delft, Netherlands.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3209281.3209413

1 OVERVIEW

Bill discussions at the California legislature contain a mixture of many interests: corporations who stand to financially benefit from the bill's passage or defeat, citizens personally impacted by the outcome of the vote, even the legislator whose reputation may be diminished for lacking the rhetorical might necessary to pass legislation promised to constituents. Beyond the opinions toward the bill itself, however, attitudes toward related entities and ideologies can often be inferred from a speaker's position on the bill[1]. For example, supporters of a bill to permit an oil pipeline can safely be considered to support the interests of oil companies; likewise, its opponents are likely to be advocates of stricter environmental regulations and the agencies that lobby for them. If the position of an organization toward a bill is known, it may then be possible to discern the alignments of legislators toward that organization based on the language they use about the bill and its related topics.

We use an existing corpus of legislative data - which includes transcripts of bill discussions, bill votes of legislators, and bill positions taken by organizations - to predict alignment scores between legislators and organizations. Table 1 contains a brief summary of this data set. We primarily focus on discussion transcripts for our predictions, as they are both a unique resource not found elsewhere and potentially contain the most sophisticated information regarding alignment. We extract relevant features from our corpus, normalizing our transcripts using standard text preprocessing techniques, and use a classification algorithm to predict the alignment between each pair of legislators and organizations.

Category	Count
Legislators	120
Committees	48
Organizations	23
Bills	1,072
Bill Discussions	1,520
Utterances	38,637

Table 1: Data Set Summary

2 EXPERIMENTS

We consider three experiments to predict alignment scores between legislators and organizations. With each method, we aggregate our features by legislator; in two of these experiments, however, we perform an additional level of aggregation: by committee or by discussion. The features we use vary slightly depending on their applicability with the experiment. The methodology and rationale for each of these experiments is discussed in the following sections.

2.1 Universal Features

The following list summarizes the features that will be used in all experiments, as detailed in Section **??**. Additional features used for specific experiments are indicated in Section **??**.

[1] https://www.digitaldemocracy.org/about

dg.o '18, May 30-June 1, 2018, Delft, Netherlands

ACM ISBN 978-1-4503-6526-0/18/05.
https://doi.org/10.1145/3209281.3209413

- **Discussion Text**: Transcribed utterances spoken by each legislator
- **Utterance Frequency**: Average number of utterances spoken by a legislator in a discussion
- **Utterance Duration**: Average number of seconds spoken by a legislator in a discussion
- **Bias Corpus Hit Rate**: Count of any word from the bias lexicon occurring in an utterance [3]
- **Sentiment Score**: Compound sentiment score from VADER, which aggregates the positive, negative, and objective scores from its sentence analysis [2]
- **Donations**: Sum of donations, gifts, and behests
- **Political Party**: Partisan affiliation of each legislator

In our experiments, we run tests on the "power set" of these features; in other words, we run tests to predict alignment scores with each feature individually, as well as every possible combination of each feature with the others.

2.2 Baselines

For each experiment, we add two baselines to the feature pool to help assess the quality of our predictions. Both of these baselines rely on a dummy classifier that simply uses the most frequent alignment category (e.g. "support") found in its training set. When the alignment categories are uniformly distributed, these dummy classifiers perform as well as random chance. However, as we know that our label distribution is not uniform, the predictions by these dummy classifiers achieve a higher accuracy than chance.

3 RESULTS

For our experiments, we group our data by different attributes, as described in Section 2; we present the results of these experiments in the following order:

(1) By Legislator
(2) By Legislator and Committee
(3) By Legislator and Discussion

In thes following plots, the x-axis indicates how many features were used for a test while the y-axis displays the average accuracy for tests with that number of features.

Because the text baseline ("__base_text__") and text feature ("__true_text__") are combined with other features for each test, their accuracy varies depending the number of features used per test. For example, for tests in which the text feature was used with four other features (i.e. five features total), the average accuracy is approximately 56%. Similarly, for the test in which only the text baseline was used, the accuracy is approximately 53.5%. Each line thus represents all tests for which that specific feature was present, with the average accuracy of a feature changing from left to right as the number of other features tested with it increases.

In our first experiment, we find that our text feature does not aid in alignment score prediction, individually performing below the overall baseline and consistently performing below the text baseline across all tests. The party feature provides the most predictive power, reaching nearly 60% accuracy individually. In our second experiment, we also see that the text feature outperforms the text baseline, as well as all but the party feature. This experiment suggests that our text feature does help predicting alignment if the samples are aggregated by legislator and committee. However, even when combined with all other features, the text feature performs only equally well as the party feature on its own, thus never improving on its accuracy.

Feature Label	Description	Experiments	Section
__freq__	Overall Baseline	ALL	2.2
__base_text__	Text Baseline	ALL	2.2
__true_text__	Text Features	ALL	??
party	Political Party	ALL	??
donated	Total Donations Received	1, 2	??
acc_donated	Donations Received to Date	3	??
count	Average Utterance Frequency	ALL	??
duration	Average Utterance Duration	ALL	??
bias	Bias Corpus Hit Rate	ALL	??
sentiment	VADER Sentiment Score	ALL	??
liberal	Committee Liberalness	2, 3	??
has_lob	Organization Lobbyist Presence	3	??
unanimity	Bill Vote Unanimity	3	??

Table 2: Feature Labels Used in Plots

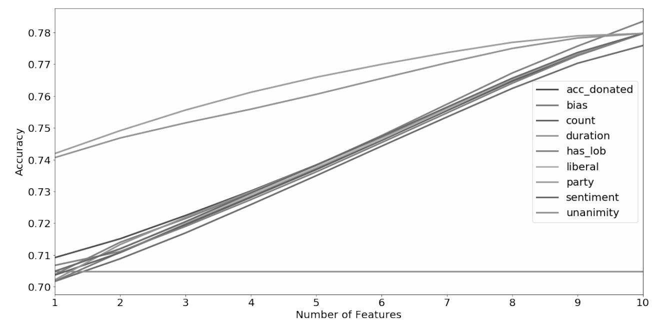

Figure 1: Experiment 3 Accuracy by Number of Features

In our final experiment, in which we use the most features, we find that, while the party feature is still dominant, the unanimity feature performs comparably well, as shown in Figure 1. We hypothesize that, since the Democratic party held a majority during the 2015-2016 session, a unanimous vote is more likely to be in agreement with organizations that support liberal causes, making agreements with unanimous votes easier to predict.

As expected, the accuracies achieved for this experiment were much higher than those of the previous experiments, reaching 78% when all features are combined. As mentioned in Section **??**, this difference is almost certainly attributed to the fact that we were predicting a binary value (agree or not) instead of ternary value (support, neutral, or oppose).

REFERENCES

[1] Pranav Anand, Marilyn Walker, Rob Abbott, Jean E Fox Tree, Robeson Bowmani, and Michael Minor. 2011. Cats rule and dogs drool!: Classifying stance in online debate. In *Proceedings of the 2nd workshop on computational approaches to subjectivity and sentiment analysis*. Association for Computational Linguistics, 1–9.

[2] Clayton J Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Eighth international AAAI conference on weblogs and social media*.

[3] Marta Recasens, Cristian Danescu-Niculescu-Mizil, and Dan Jurafsky. 2013. Linguistic Models for Analyzing and Detecting Biased Language.. In *ACL (1)*. 1650–1659.