

Analysis of Nuanced Stances and Sentiment Towards Entities of US Politicians through the Lens of Moral Foundation Theory

Shamik Roy

Department of Computer Science
Purdue University, USA
roy98@purdue.edu

Dan Goldwasser

Department of Computer Science
Purdue University, USA
dgoldwas@purdue.edu

Abstract

The Moral Foundation Theory suggests five moral foundations that can capture the view of a user on a particular issue. It is widely used to identify sentence-level sentiment. In this paper, we study the nuanced stances and partisan sentiment towards entities of US politicians using Moral Foundation Theory, on two politically divisive issues - *Gun Control* and *Immigration*. We define the nuanced stances of the US politicians on these two topics by the grades given by related organizations to the politicians. To conduct this study, we first filter out 74k and 87k tweets on the topics *Gun Control* and *Immigration*, respectively, from an existing tweet corpus authored by US parliament members. Then, we identify moral foundations in these tweets using deep relational learning. Finally, doing qualitative and quantitative evaluations on this dataset, we found out that there is a strong correlation between moral foundation usage and politicians' nuanced stances on a particular topic. We also found notable differences in moral foundation usage by different political parties when they address different entities.

1 Introduction

Over the last decade political discourse has shifted from traditional news outlet to social media. These platforms give politicians the means to interact with their supporters and explain their political perspectives and policy decisions. While formulating policies and passing legislation are complex processes which require reasoning over the pros and cons of different alternatives, *gathering support* for these policies often relies on appealing to peoples' "gut feeling" and invoking an emotional response (Haidt, 2001).

Moral Foundation Theory (MFT) provides a theoretical framework for analyzing the use of moral sentiment in text. The theory (Haidt and Joseph, 2004; Haidt and Graham, 2007) suggests that there

are a small number of moral values, emerging from evolutionary, cultural and social reasons, which humans support. These are referred to as the moral foundations (MF) and include Care/Harm, Fairness/Cheating, Loyalty/Betrayal, Authority/Subversion, and Purity/Degradation. This theory was used to explain differences between political ideologies, as each side places more or less value on different moral foundations (Graham et al., 2009). Liberals tend to emphasize the *Fairness* moral foundation, for example, consider the following tweet discussing the 2021 mass shooting event in Colorado, focusing on how the race of the shooter changes the coverage of the event.

Liberal Gun Control tweet. Fairness

@IlhanMN The shooter's race or ethnicity seems front and center when they aren't white. Otherwise, it's just a mentally ill young man having a bad day.

On the other hand, conservatives tend to place more value on *Loyalty*. The following tweet discusses the same event, emphasizing solidarity with the families of victims and the broader community.

Conservative Gun Control tweet. Loyalty

@RepKenBuck My prayers are with the families of the victims of today's tragedy in Boulder. I join the entire community of Boulder in grieving the senseless loss of life. I am grateful for the officers who responded to the scene within minutes. You are true heroes.

In this paper, we study the relationship between moral foundation usage by politicians on social media and the stances they take on two policy issues, *Gun Control* and *Immigration*. We use the dataset provided by (Johnson and Goldwasser, 2018) for training a model for automatically identifying moral foundations in tweets. We then apply the model to a collection of 74k and 87k congressional tweets discussing the two issues - *Gun Control* and *Immigration*, respectively. Our analysis goes beyond binary liberal-conservative ideological labels (Preotiuc-Pietro et al., 2017). We use a scale of 5 letter grades assigned to politicians by

relevant policy watchdog groups, based on their votes on legislation pertaining to the specific policy issue. We analyze the tweets associated with the members of each group. Furthermore, we hypothesize that even when different groups use similar moral foundation, they aim to invoke different feelings in the readers. To capture these differences, we analyze the targets of the moral tweets by different groups. Our analysis captures several interesting trends. First, the proportion of *non-moral* tweets on both issues decreases as grades move from A (most conservative) to F (most liberal), while for the topic of Gun Control (Immigration), the proportion of *Harm (Loyalty)* tweets increases. Second, even when using the same moral foundation, their targets differ. For example, when discussing Gun Control, using the *Loyalty* moral foundation, liberal mostly mention *march life, Gabby Gifford*, while conservatives mention *gun owner, Texas*.

2 Related Works

The Moral Foundation Theory (MFT) (Haidt and Joseph, 2004; Haidt and Graham, 2007) has been proven to be useful in explaining social behaviour of humans (Mooijman et al., 2018; Hoover et al., 2018; Dehghani et al., 2016; Brady et al., 2017; Hoover et al., 2020). Recent works have shown that political discourse can also be explained using MFT (Dehghani et al., 2014; Johnson and Goldwasser, 2018, 2019). Existing works explain the political discourse mostly at issue and sentence level (Fulgoni et al., 2016; Garten et al., 2016; Lin et al., 2018; Xie et al., 2019) and at left-right polar domains of politics.

Several works have looked at analyzing political ideologies, beyond the left and right divide, using text (Sim et al., 2013; Preoŕiuc-Pietro et al., 2017), and specifically using Twitter data (Conover et al., 2011; Johnson and Goldwasser, 2016; Mohammad et al., 2016; Demszky et al., 2019). To the best of our knowledge, this is the first work that studies whether MFT can be used to explain nuanced political standpoints of the US politicians, breaking the left/right political spectrum to nuanced standpoints. We also study the correlation between entity mentions and moral foundation usage by different groups, which helps pave the way to analyze partisan sentiment towards entities using MFT. In that sense, our work is broadly related to entity-centric affective analysis (Deng and Wiebe, 2015; Field and Tsvetkov, 2019; Park et al., 2020).

We use a deep structured prediction approach (Pacheco and Goldwasser, 2021) to identify moral foundations in tweets by being motivated from the works that combine structured prediction with deep neural networks in NLP tasks (Niculae et al., 2017; Han et al., 2019; Liu et al., 2019; Widmoser et al., 2021).

3 Dataset

In this section, we describe the data collection process to analyze the US politicians' stances and sentiment towards entities on the topics - *Immigration* and *Gun Control*. First, we discuss existing datasets. Then, we create a topic specific lexicon from existing resource to identify topics in new data. Finally, we collect a large tweet corpus on the two topics using a lexicon matching approach.

3.1 Candidate Datasets

To study the nuanced stances and sentiment towards entities of politicians using MFT on the text they use, ideally, we need a text dataset annotated for moral foundations from US politicians with known political bias. To the best of our knowledge there are two existing Twitter datasets that are annotated for moral foundations - (1) The Moral Foundations Twitter Corpus (MFTC) by Hoover et al. (2020), and (2) The tweets by US politicians by Johnson and Goldwasser (2018). In MFTC, the moral foundation annotation is done in 35k Tweets on 7 distinct domains, some of which are not related to politics (e.g. Hurricane Sandy) and the political affiliations of the authors of the tweets are not known. The dataset proposed by Johnson and Goldwasser (2018) contains 93K tweets by US politicians in the years 2016 and 2017. 2050 of the tweets are annotated for moral foundations, policy frames (Boydston et al., 2014) and topics. The dataset contains 6 topics including *Gun Control* and *Immigration*. We extend this dataset for these two topics by collecting more tweets from US Congress members using a lexicon matching approach, described in the next section.

3.2 Building Topic Indicator Lexicon

To build a topic indicator lexicon, we take the dataset proposed by Johnson and Goldwasser (2018). We build topic indicator lexicons for each of the 6 topics comprised of n-grams ($n \leq 5$) using Pointwise Mutual Information (PMI) scores (Church and Hanks, 1990). For an n-gram, w we

calculate the pointwise mutual information (PMI) with topic t , $I(w, t)$ using the following formula.

$$I(w, t) = \log \frac{P(w|t)}{P(w)}$$

Where $P(w|t)$ is computed by taking all tweets with topic t and computing $\frac{\text{count}(w)}{\text{count}(\text{all}n\text{grams})}$ and similarly, $P(w)$ is computed by counting n-gram w over the set of tweets with any topic. Now, we rank n-grams for each topic based on their PMI scores. We assign one n-gram to its highest PMI topic only. Then for each topic we manually go through the n-gram lexicon and omit any n-gram that is not related to the topic. In this manner, we found an indicator lexicon for each topic. The lexicons for the topics *Gun Control* and *Immigration* can be found in Appendix A. Note that, as a pre-processing step, n-grams were stemmed and singularized.

3.3 Tweet Collection

We use the large number of unlabeled tweets from US Congress members, written between 2017 and February, 2021¹. We detect tweets related to the topics *Gun Control* and *Immigration* using lexicon matching. If a tweet contains any n-gram from the topic lexicons, we label the tweet with the corresponding topic. We take only the tweets on topics *Gun Control* and *Immigration* from the Democrat and Republican US Congress members for our study. Given the political affiliation of the authors of the tweets, this dataset is readily useful for the analysis of political stance and partisan sentiment. The details of the dataset is presented in Table 1.

| | GUN CONTROL | | | IMMIGRATION | | |
|-------------------|-------------|--------|--------|-------------|--------|--------|
| | DEM | REP | TOTAL | DEM | REP | TOTAL |
| # of politicians | 350 | 377 | 727 | 349 | 364 | 713 |
| # of Twitter acc. | 644 | 641 | 1,285 | 621 | 606 | 1,227 |
| # of tweets | 53,793 | 20,424 | 74,217 | 65,671 | 21,407 | 87,078 |

Table 1: Dataset summary. Here, ‘Dem’ and ‘Rep’ represent ‘Democrat’ and ‘Republican’, respectively. The number of politicians and the number of Twitter accounts differs as politicians often have multiple accounts (e.g. personal account, campaign account, etc.).

4 Identification of Moral Foundation in Tweets

To identify moral foundations in the collected dataset, we rely on a supervised approach using

¹<https://github.com/alexlitel/congressstweets>

a deep relational learning framework. In this section, we first describe the model we use for the supervised classification. Then, we describe our training procedure and analyze the performance of our model on a held out set. Finally, we describe the procedure to infer moral foundations in the collected dataset using our model.

4.1 Deep Relational Learning For Moral Foundation (MF) Identification

For the identification of moral foundation (MF) in tweets, [Johnson and Goldwasser \(2018\)](#) rely on linguistic cues such as - political slogans, Policy Frames, Annotator’s Rationale; along with party affiliation, topic and so on, while [Johnson and Goldwasser \(2019\)](#) models the behavioural aspects of the politicians in MF identification. In both of the works they use Probabilistic Soft Logic for modeling. Some of the features used by [Johnson and Goldwasser \(2018\)](#) and [Johnson and Goldwasser \(2019\)](#) are hard to get for a large corpus and some require human annotation. Note that, in this section, our goal is not to outperform the state-of-the-art MF classification results, rather we want to identify MFs in the large corpus where only limited information is available. So, to identify MFs in our corpus we mostly rely on text and the information available with the unlabeled corpus such as, topics, authors’ political affiliations and time of the tweets. We jointly model all of these features using DRaiL, a declarative framework for deep structured prediction proposed by [Pacheco and Goldwasser \(2021\)](#) which is described below.

Modeling Features and Dependencies In DRaiL, we can explicitly model features such as - tweet text, authors’ political affiliations and topics using *base rules* as follows.

$$\begin{aligned} r_1 &: \text{Tweet}(t) \Rightarrow \text{HasMF}(t, m) \\ r_2 &: \text{Tweet}(t) \wedge \text{HasIdeology}(t, i) \Rightarrow \text{HasMF}(t, m) \\ r_2 &: \text{Tweet}(t) \wedge \text{HasTopic}(t, k) \Rightarrow \text{HasMF}(t, m) \end{aligned}$$

These rules correspond to base classifiers that map features in the left hand side of the \Rightarrow to the predicted output in the right hand side. For example, the rule, r_2 translates as "A tweet, t with authors’ political affiliation, i has moral foundation label, m ". We can also model the temporal dependency between two classification decisions using a second kind of rule, namely *constraint* as follows.

$$\begin{aligned} c &: \text{SameIdeology}(t_1, t_2) \wedge \text{SameTopic}(t_1, t_2) \wedge \\ &\text{SameTime}(t_1, t_2) \wedge \text{HasMF}(t_1, m) \Rightarrow \text{HasMF}(t_2, m) \end{aligned}$$

This constraint translates as "If two tweets have the same topic, are from the authors of the same political affiliation and are published nearly at the same time, then they have the same moral foundation". This constraint is inspired from the experiments done by Johnson and Goldwasser (2019). In DRaiL, rules can be weighted or unweighted. We consider weighted version of the rules, making constraint c a soft-constraint as it is not guaranteed to be true all of the time. In DRaiL, the global decision is made considering all rules. It transforms rules into linear inequalities and MAP inference is then defined as an integer linear program:

$$\begin{aligned} \mathbf{y} \in \{0,1\}^n P(\mathbf{y}|\mathbf{x}) \equiv \mathbf{y} \in \{0,1\}^n \sum_{\psi_r, t \in \Psi} w_r \psi_r(\mathbf{x}_r, \mathbf{y}_r) \\ \text{s.t. } c(\mathbf{x}_c, \mathbf{y}_c) \leq 0; \quad \forall c \in C \end{aligned} \quad (1)$$

Here, rule grounding, r , generated from template, t , with input features, \mathbf{x}_r and predicted variables, \mathbf{y}_r defines the potential, $\psi_r(\mathbf{x}_r, \mathbf{y}_r)$ where weights, w_r are learned using neural networks defined over parameter set, θ . The parameters can be learned by training each rule individually (*locally*), or by using inference to ensure that the scoring functions for all rules result in a globally consistent decision (*globally*) using the structured hinge loss:

$$\max_{\hat{\mathbf{y}} \in Y} (\Delta(\hat{\mathbf{y}}, \mathbf{y}) + \sum_{\psi_r \in \Psi} \Phi_t(\mathbf{x}_r, \hat{\mathbf{y}}_r; \theta^t)) - \sum_{\psi_r \in \Psi} \Phi_t(\mathbf{x}_r, \mathbf{y}_r; \theta^t)$$

Here, t is rule template, Φ_t is the associated neural network, and θ^t is the parameter set. \mathbf{y} and $\hat{\mathbf{y}}$ are gold assignments and predictions resulting from the MAP inference, respectively.

Neural Architectures Each base rule and the soft-constraint is associated with a neural architecture which serve as weighting functions for the rules and constraints. For rules, r_1 , r_2 and r_3 , we use BERT (Devlin et al., 2019) to encode the tweet text. In rules r_2 and r_3 , we encode ideology and topic with a feed-forward neural network over their one-hot encoded form and we concatenate the encoded features with BERT representation of tweets to get a final representation for the rule. In all of the rules we use a classifier on top of the final representation that maps the features to labels. For the soft-constraint c , we encode the ideologies and topics in the left hand side of the constraint similarly and concatenate them and pass through a classifier to predict if the constraint holds or not.

4.2 Experimental Evaluation

We use the dataset proposed by Johnson and Goldwasser (2018) for this experiment.² We perform a 5-fold cross validation on 2050 tweets annotated for moral foundations. This is a 11 class classification task where there is one additional class, ‘Non-moral’ apart from the 10 moral classes. We experiment with the global learning of DRaiL using rules r_1, r_2, r_3 and soft constraint c . For the BERT (base-uncased) classifiers we use a learning rate of $2e^{-5}$, batch size of 32, patience 10 and AdamW as optimizer. All of the tweets were truncated to a length of 100 tokens before passing through BERT. For constraint c we consider two tweets to be at the same time if they are published on the same day. All of the one-hot representations are mapped to a 100 dimensional space and ReLU and Softmax activation functions are used in all hidden and output neural units, respectively. The hyper-parameters are determined empirically.³ We compare our model with two baselines as follows.

(1) Lexicon matching with Moral Foundations

Dictionary (MFD) This approach does not have a training phase. Rather we use the Moral Foundation Dictionary (Graham et al., 2009) and identify moral foundation in a tweet using unigram matching from the MFD. A tweet having no dictionary matching is labeled as ‘Non-moral’.

(2) Bidirectional-LSTM

We run a bidirectional LSTM (Hochreiter and Schmidhuber, 1997) over the Glove (Pennington et al., 2014) word embeddings of the words of the tweets. We concatenate the hidden states of the two opposite directional LSTMs to get representation over one timestamp and average the representations of all time stamps to get the final representation of a tweet. We map each tweet to a 128-d space using Bi-LSTM and use this representation for moral foundation classification using a fully connected output layer. We use the same folds as the DRaiL experiments.

The classification results are summarized in Table 2. We can see that the DRaiL model combining all base rules and the soft-constraint performs best. This indicates that combining other features with

²More details on dataset can be found in the original paper.

³Dataset and codes can be found at <https://github.com/ShamikRoy/MF-Prediction>.

| | MODELS | AVERAGE MACRO F1 | WEIGHTED F1 |
|-----------|-----------------------|---------------------|----------------|
| Baselines | MFD matching | 15.93 | 18.38 |
| | Bi-LSTM | 42.59 | 50.43 |
| DRaiL | r_1 (BERT only) | 49.01 | 57.96 |
| | $r_1 + r_2$ | 50.54 | 58.90 |
| | $r_1 + r_2 + r_3$ | 51.49 | 60.02 |
| | $r_1 + r_2 + r_3 + c$ | 52.14 | 60.24 |

Table 2: Moral Foundation classification results.

BERT and modeling the dependencies among multiple decisions help in prediction. This encourages us to experiment with other linguistic features (e.g. policy frames) and dependencies as a future work.

| MORALS | PREC. | REC. | F1 | SUPPORT |
|-------------|-------|-------|-------|---------|
| CARE | 53.18 | 62.02 | 57.26 | 337 |
| HARM | 52.01 | 56.35 | 54.10 | 252 |
| FAIRNESS | 67.93 | 59.24 | 63.29 | 211 |
| CHEATING | 27.27 | 16.98 | 20.93 | 53 |
| LOYALTY | 52.63 | 56.60 | 54.55 | 212 |
| BETRAYAL | 60.00 | 31.58 | 41.38 | 19 |
| AUTHORITY | 40.17 | 41.59 | 40.87 | 113 |
| SUBVERSION | 68.55 | 71.23 | 69.86 | 358 |
| PURITY | 67.20 | 64.62 | 65.88 | 130 |
| DEGRADATION | 53.85 | 22.58 | 31.82 | 31 |
| NON-MORAL | 77.48 | 70.06 | 73.58 | 334 |
| ACCURACY | | | 60.39 | 2050 |
| AVG. MAC. | 56.39 | 50.26 | 52.14 | 2050 |
| WEIGHTED | 60.72 | 60.39 | 60.24 | 2050 |

Table 3: Per class Moral Foundation classification results for the best model in Table 2

We present the per class statistics of the prediction of the best model in Table 3. We can see that mostly the classes with lower number of examples are harder to classify for the model (e.g. *Cheating*, *Degradation*). So, annotating more tweets on the low frequency classes may improve the overall performance of the model.

4.3 Inference on the Collected Corpus

Now, we train our best model (combining all base rules and the constraint in DRaiL) using the dataset we experiment with in Section 4.2. We held out 10% of the data as validation set selected by the random seed of 42. We train the model using the hyper-parameters described in Section 4.2 and predict moral foundations in the tweets of the large corpus we annotated for the topics *Gun Control* and *Immigration* in Section 3. The validation macro F1 score and weighted F1 scores of the model were 49.44% and 58.30%, respectively. We use this annotated dataset to study nuanced stances and partisan sentiment towards entities of the US politicians.

5 Analysis of Politicians’ Nuanced Stances

In this section, we analyze the nuanced stances of US politicians on the topics *Gun Control* and *Immigration*, using Moral Foundation Theory. First, we define nuanced political stances. Then we study the correlation between the moral foundation usage and nuanced political stances.

5.1 Nuanced Political Stance

Despite of being highly polarized, US politicians show mixed stances on different topics. For example, a politician may be supportive of gun prevention laws to some extent despite their party affiliation of the Republican Party. So, we hypothesize that the political stance is more nuanced than binary, left and right. We define the nuanced political stances of the politicians as the grades assigned to them by the *National Rifle Association (NRA)*⁴ on *Gun Control* and by *NumbersUSA*⁵ on *Immigration*. The politicians are graded in range (A+, A, . . . , F, F-) based on candidate questionnaire and their voting records by both of the organizations in the two different topics where A+ indicates most anti-immigration/pro-gun and F or F- indicates the most pro-immigration/anti-gun. In other words, A+ means extreme right and F/F- means extreme left and the other grades fall in between. We convert these letter grades in 5 categories: A, B, C, D, F. Here, A+, A and A- grades are combined in A and so on. We define these grades as nuanced stances of the politicians on the two topics.

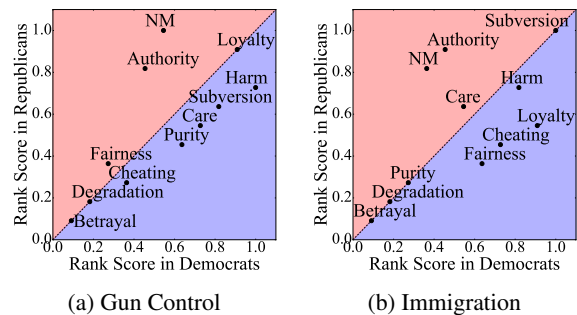


Figure 1: Polarization in Moral Foundation usage. Here, NM stands for ‘Non-moral’.

5.2 Moral Foundation Usage

In this section, first, we study the political polarization, similar to [Roy and Goldwasser \(2020\)](#), in

⁴Collected from [everytown.org](#)

⁵Collected from [numbersusa.com](#)

moral foundation usage by Democrats and Republicans on the two topics. Therefore, we rank the moral foundations by the frequency of usage inside each party. Then we plot the rank score of each moral foundation in Democrats and Republicans in x and y axes, respectively, where the most used moral foundation gets the highest rank score. Any moral foundation falling in the diagonal is not polarized and as far we go away from the diagonal it becomes more polarized. We show the polarization graphs for the two topics in Figure 1. It can be seen that the parties are polarized in moral foundation usage. The Republicans use ‘Non-moral’ and ‘Authority’ moral foundations in both of the topics. On the other hand, Democrats use ‘Subversion’ and ‘Harm’ on *Gun Control* and ‘Loyalty’ and ‘Cheating’ on *Immigration*.

Now, we examine the moral foundation usage by the politicians from each of the grade categories. For that, we match the politicians with grades with our dataset and consider politicians tweeting at least 100 times on each topic. The statistics of politicians and corresponding tweets found for each grade is presented in Table 4. Now, to compare


| GRADES | GUN CONTROL | | IMMIGRATION | |
|--------|---------------|----------|---------------|----------|
| | # POLITICIANS | # TWEETS | # POLITICIANS | # TWEETS |
| A | 31 | 6,822 | 25 | 5,592 |
| B | 5 | 1,236 | 11 | 2,177 |
| C | 7 | 908 | 3 | 679 |
| D | 9 | 1,340 | 14 | 4,691 |
| F | 128 | 33,792 | 123 | 38,102 |

Table 4: Distribution of number of Politicians and tweets over the letter grades.


the moral foundation usage by each of the grade classes, we rank the moral foundations based on their usage inside each grade. Then we compare the rank of each grade class with the two opposite extremes (grades A and F) using Spearman’s Rank Correlation Coefficient (Zar, 2005) where coefficient 1 means perfect correlation. As the grades B, C, D have fewer tweets, we sub-sample 500 tweets from each class and do the analysis on them. We repeat this process 10 times with 10 different random seeds and plot the average correlations in Figure 2.⁶

It can be seen from the figures that the correlations follow a progressive trend with the extreme left while moving from grade A to grade F and the trend is opposite with the extreme right, for both of the topics. This indicates that there is a correlation

⁶Standard Deviations can be found in Appendix B.




(a) Gun Control




(b) Immigration


Figure 2: Correlation of moral foundation usage with NRA and NumbersUSA grades of politicians on the topics *Gun Control* and *Immigration*, respectively.




(a) Usage of ‘Non-moral’ on *Gun Control*



(b) Usage of ‘Harm’ on *Gun Control*



(c) Usage of ‘Non-moral’ on *Immigration*



(d) Usage of ‘Loyalty’ on *Immigration*

Figure 3: Moral Foundation distribution over politicians’ grades.

between the MF usage and politicians’ nuanced stances. To further analyze which moral foundations most correlate with the nuanced stances, we plot the percentage of usage of the most polar moral foundations from Figure 1, inside each grade class. We found good correlations in case of the usage of ‘Non-moral’ and ‘Harm’ on *Gun Control*; in usage of ‘Non-moral’ and ‘Loyalty’ on *Immigration*. The distributions are shown in Figure 3. Distribution plots for the other moral foundations can be found

| | GUN CONTROL | | IMMIGRATION | |
|-------------|---|---|--|--|
| Morals | High PMI Entities by Democrats | High PMI Entities by Republicans | High PMI Entities by Democrats | High PMI Entities by Republicans |
| Care | community safe, gun violence prevention, assault weapon | law enforcement, bipartisan bill, health care | protect dreamer, immigration status, young people | build wall, immigration law, border patrol |
| Harm | mass shooting, innocent life, school shooting | police officer, mexico, texas | detention facility, detention center, migrant child | illegal alien, build wall, illegal immigrant |
| Fairness | gun sale, universal background check, gun owner | gun owner, amendment, nra | immigration status, dream promise, dream | illegal immigrant, illegal alien, american citizen |
| Cheating | bump stock, nra, black | gun owner, gun control, amendment | citizenship question, muslim, american value | illegal immigrant, illegal alien, illegal immigration |
| Loyalty | march life, gabby gifford, young people | gun owner, texas, charleston | protect dream, defend daca, dream promise act | border patrol, southern border, american people |
| Betrayal | congress, gun | gun | human right, refugee, american citizen | illegal alien, illegal immigrant, sanctuary city |
| Authority | bipartisan background check, american people, house judiciary | gun, american people | circuit judge, comprehensive immigration reform, supreme court | circuit judge, circuit court, senate |
| Subversion | house gop, republican, gun lobby | gun control, dem, medicare | trump shutdown, national emergency, border wall | illegal immigrant, illegal immigration, sanctuary city |
| Purity | pulse, tragic shooting, honor action | tragic shooting, police officer, las vegas | refugee, america, american value | american citizen, circuit court, illegal alien |
| Degradation | el paso, nra, republican | orlando, texas, black | muslim, usc, daca | muslim, human right, fema |
| Non-moral | town hall, medicare, shannon r watt | amendment, gun, charleston | medicare, usc, house judiciary | government shutdown, border security, homeland |

Table 5: Top-3 high PMI entities for each moral foundation by each party.

in Appendix C. It can be seen from the figures that, as we move from grade A to F, the usage of ‘Non-moral’ decreases for both of the topics, indicating - the more conservative a politician is, they discuss the issues from a more ‘Non-moral’ perspective. On the other hand, more usage of ‘Harm’ and ‘Loyalty’ indicates more liberal stances on *Gun Control* and *Immigration*, respectively.

6 Analysis of Partisan Sentiment Towards Entities

In this section, we study the partisan sentiment towards entities by examining the usage of moral foundations while discussing the entities. First, we extract entities from the tweets, then we analyze the usage of moral foundations in the context of those entities by the two opposite parties.

6.1 Entity Extraction from Tweets

To study partisan sentiment towards entities we first identify entities mentioned in the tweets. We hypothesize entities to be noun phrases. So, we use an off-the-shelf noun phrase extractor⁷ and extract noun phrases from the tweets. We filter out noun phrases occurring less than 100 times. Then we

⁷<https://textblob.readthedocs.io/>

manually filter out noun phrases that are irrelevant to the topics (e.g. COVID-19). In this manner, we found 64 and 79 unique noun phrases for *Gun Control* and *Immigration*, respectively. We treat these noun phrases as entities and run our analysis using these entities. The complete list of entities can be found in Appendix D

6.2 MF Usage in the Context of Entities

In this section, we analyze the partisan sentiment towards entities by looking at the moral foundation usage trend of the parties when discussing the entities related to the topics. For each party and each moral foundation we calculate the PMI score with each entity. We create 22 classes comprised of the 2 party affiliations and 11 moral foundation classes (e.g. Democrat-Care, Republican-Care and so on) and calculate the PMI scores as described in Section 3. We list the top-3 highest PMI entities for each moral foundation and each party in Table 5. We can see notable difference in moral foundation usage in the context of different entities by the two parties. For example, on the issue *Immigration*, the Democrats use ‘Care’ when addressing ‘dreamers’ and ‘young people’. On the other hand, the Republicans use care in the context of ‘border wall’ and ‘border patrol’. On the

| ID | Party | Topic | Entities | Predicted MF | Tweet Text |
|-----|------------|-------------|---------------------------------------|--------------|---|
| (1) | Democrat | Immigration | Migrant Child; Trump Administration | Harm | How many more migrant children must die under the Trump administration until something changes? |
| (2) | Democrat | Immigration | Migrant Child; Detention Facility | Harm | 12,800! That’s how many migrant CHILDREN are locked up in detention facilities in America. How can this be happening? |
| (3) | Republican | Gun Control | Police Officer | Harm | A Charleston police officer has been shot in the face. |
| (4) | Republican | Gun Control | Police Officer; Communities; Families | Care | North Carolina police officers protect our communities , keep our families safe, and have earned our support. |
| (5) | Democrat | Gun Control | Tragic Shooting | Purity | Our thoughts and prayers are with the victims of the tragic shooting in Las Vegas. Look forward to a full investigation to give us answers. |
| (6) | Republican | Gun Control | Tragic Shooting | Purity | Praying for the families of victims of tragic shooting in Vegas. Time to transcend politics and pray for God’s healing for those affected. |
| (7) | Republican | Gun Control | Gun Owner | Fairness | Law-abiding gun owners deserve the full protection of the U.S. Constitution when exercising their right to carry a concealed weapon – and that right should not end at a state line. |
| (8) | Democrat | Gun Control | Gun Owner | Fairness | I am a hunter who believes in protecting the rights of law abiding gun owners . I am also a father of two young boys who believes there need to be changes in our gun laws. |

Table 6: Qualitative evaluation of Moral Foundation usage in the context of entities.

issue *Gun Control*, when talking about ‘NRA’ the Democrats associate ‘Cheating’ and ‘Degradation’, while the Republicans use ‘Fairness’. These imply high polarization in partisan sentiment towards entities. We can see some interesting cases as well. For example, on *Guns*, the Republicans use ‘Harm’ with the entity ‘police officer’ and on *Immigration*, the Democrats use ‘Harm’ with ‘migrant child’. On *Guns*, democrats and republicans sometimes use the same moral foundation in the context of the same entity. For example, both Democrats and Republicans use ‘Fairness’ in the context of ‘Gun Owner’ and ‘Purity’ in the context of ‘tragic shooting’. So, we take a closer look at the usage of MFs in the context of these entities and list a few tweets discussing each of these entities in Table 6.

We can see that on *Immigration*, for Democrats, ‘migrant child’ is target of harm while ‘detention facility’ and ‘Trump administration’ are the entities posing the harm (examples (1), (2) in Table 6). So, even if the high-level moral foundation is the same, different participating entities in the text may have different partisan sentiments towards them.

On *Guns*, although the entity ‘police officer’ carries a positive sentiment for the Republicans across different moral foundations, the fine-grained sentiment towards this entity is different in the case of different moral foundations. For example, ‘police officer’ is the target of harm and is the entity providing care for the Republicans when used in

the context of ‘Harm’ and ‘Care’, respectively (examples (3), (4) in Table 6). So, moral foundation can explain the sentiment towards entities beyond positive and negative categories.

In the context of ‘Gun Owner’, both of the parties use ‘Fairness’ in support of gun owners’ rights, but they frame the issue differently - Democrats, by focusing on the need for more restrictions while preserving gun rights (example (8)) and Republicans, by focusing on the violation of constitutional rights if more restrictions are applied (example (7)). So, even if the moral foundation usage is the same, there is a framing effect to establish the corresponding partisan stances. While using ‘Purity’ in the context of ‘tragic shooting’, we found that both of the parties express their prayers for the shooting victims (example (5), (6)).

Now, we find out the entities with highest disagreement between parties in moral foundation usage in context. To calculate the disagreement we rank the moral foundations based on frequency in usage by each party in the context of each entity. Then we calculate the Spearman’s Rank Correlation Coefficient between these two rankings for each entity and list the top-10 entities with the highest disagreement in Table 7. Then we show the polarity graphs for one entity from each topic list in Figure 4. We can see that, on *Gun*, while discussing ‘Amendment’ the Republicans use ‘Loyalty’, although ‘Loyalty’ is not polarized towards the Re-

publicans in aggregate (Figure 1). On the other hand, the Democrats use ‘Cheating’ in the context of ‘Amendment’. Similarly, while discussing ‘Donald Trump’ on *Immigration*, the Democrats use ‘Cheating’ more, while the Republicans use ‘Care’ and ‘Authority’. These analyses indicate that moral foundation analysis can be a useful tool to analyze partisan sentiment towards entities.

| TOPICS | ENTITIES WITH HIGHEST DISAGREEMENT IN MF USAGE IN CONTEXT BETWEEN DEMOCRATS AND REPUBLICANS |
|--------------------|--|
| GUN CONTROL | Amendment, background check, gun, gun control, NRA, gun violence, violence, Congress, gun owner, high school |
| IMMIGRATION | immigration policy, Donald Trump, America, DHS Gov, Supreme Court, legal immigration, Mexico, immigration system, DHS, ICE |

Table 7: Top-10 entities with highest disagreement in MF usage in context between Democrats and Republicans (in descending order of agreement).

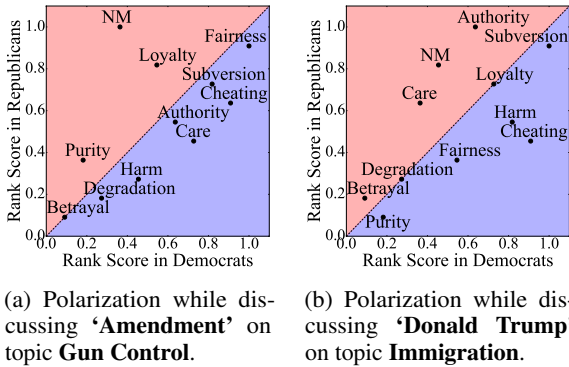


Figure 4: Polarization in entity discussion.

7 Future Work

In this section, we discuss some potential research directions that our analyses may lead to and their application in understanding political discourse.

Our experiments in Section 4 show that joint modeling of multiple aspects of the dataset (e.g. text, issue, and political affiliation) and the dependency among multiple decisions (e.g. temporal dependency), helps in classification. Incorporating other information such as linguistic cues, behavioural aspects, and so on, has the potential to improve the prediction furthermore. In general, incorporating information from multiple sources (e.g. social, textual) and modeling dependencies among decisions is an interesting future work that can help

in the identification of the underlying intent of the text. So, this framework may be extended to similar tasks, such as political framing analysis, misinformation analysis, propaganda detection, and so on.

In Section 5, we found out that moral foundation usage can be useful in explaining the nuanced political stances of politicians beyond the left/right discreet categories. We observed that usage of some moral foundations strongly correlates with the nuanced stances of the politicians. While the stances of the extreme left (grade F) and extreme right (grade A) politicians are easy to explain, what are the stances of the politicians in the middle (grades B to D), is yet to be investigated qualitatively. This line of research would help in understanding the stance of the politicians at **individual** levels and has real-life implications. For example, understanding politicians’ individual stances would help determine their future vote on legislative decisions and to identify the aisle-crossing politicians.

In Section 6, we found out clear cases where sentiment towards entities can be explained by grounding the Moral Foundation Theory at the entity level. This is an interesting direction where we can seek answers to several research questions, such as, (r1) What are the dimensions in a moral foundation category along which the sentiment towards the entities can be explained?; (r2) Can sentiment towards entities, inspired from moral foundations, explain political discourse?; (r3) Do the sentiment towards entities change over time and in response to real-life events? We believe our analyses will help advance the research in this direction.

8 Summary

In this paper, we study how Moral Foundation Theory (MFT) can explain nuanced political stances of US politicians and take the first step towards partisan sentiment analysis targeting different entities using MFT. We collect a dataset of 161k tweets authored by US politicians, on two politically divisive issues, *Gun Control* and *Immigration*. We use a deep relational learning approach to predict the moral foundations in the tweets, that models tweet text, topic, author’s ideology, and captures temporal dependencies based on publication time. Finally, we analyze the politicians’ nuanced standpoints and partisan sentiment towards entities using MFT. Our analyses show that both phenomena can be explained well using MFT, which we hope will help motivate further research in this area.

9 Ethical Considerations

To the best of our knowledge no code of ethics was violated throughout the experiments and data collection done in this paper. We presented the detailed data collection procedure and cited relevant papers and websites from which we collected the data. We provided all implementation details and hyper-parameter settings for reproducibility. Any qualitative result we report is outcome from machine learning models and doesn't represent the authors' personal views, nor the official stances of the political parties analyzed.

Acknowledgements

We gratefully acknowledge Maria Leonor Pacheco for helping in setting up the deep relational learning task using DRaiL and the anonymous reviewers for their insightful comments. We also acknowledge Nikhil Mehta for his useful feedback on the writing.

References

- Amber Boydston, Dallas Card, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2014. Tracking the development of media frames within and across policy issues.
- William J Brady, Julian A Wills, John T Jost, Joshua A Tucker, and Jay J Van Bavel. 2017. Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, 114(28):7313–7318.
- Kenneth Ward Church and Patrick Hanks. 1990. [Word association norms, mutual information, and lexicography](#). *Computational Linguistics*, 16(1):22–29.
- 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.
- Morteza Dehghani, Kate Johnson, Joe Hoover, Eyal Sagi, Justin Garten, Niki Jitendra Parmar, Stephen Vaisey, Rumen Iliev, and Jesse Graham. 2016. Purity homophily in social networks. *Journal of Experimental Psychology: General*, 145(3):366.
- Morteza Dehghani, Kenji Sagae, Sonya Sachdeva, and Jonathan Gratch. 2014. Analyzing political rhetoric in conservative and liberal weblogs related to the construction of the “ground zero mosque”. *Journal of Information Technology & Politics*, 11(1):1–14.
- Dorottya Demszky, Nikhil Garg, Rob Voigt, James Zou, Jesse Shapiro, Matthew Gentzkow, and Dan Jurafsky. 2019. Analyzing polarization in social media: Method and application to tweets on 21 mass shootings. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2970–3005.
- Lingjia Deng and Janyce Wiebe. 2015. [Joint prediction for entity/event-level sentiment analysis using probabilistic soft logic models](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 179–189, Lisbon, Portugal. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Anjalie Field and Yulia Tsvetkov. 2019. Entity-centric contextual affective analysis. *arXiv preprint arXiv:1906.01762*.
- Dean Fulgoni, Jordan Carpenter, Lyle Ungar, and Daniel Preoțiuc-Pietro. 2016. [An empirical exploration of moral foundations theory in partisan news sources](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3730–3736, Portorož, Slovenia. European Language Resources Association (ELRA).
- Justin Garten, Reihane Boghrati, Joe Hoover, Kate M Johnson, and Morteza Dehghani. 2016. Morality between the lines: Detecting moral sentiment in text. In *Proceedings of IJCAI 2016 workshop on Computational Modeling of Attitudes*.
- Jesse Graham, Jonathan Haidt, and Brian A Nosek. 2009. Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology*, 96(5):1029.
- Jonathan Haidt. 2001. The emotional dog and its rational tail: a social intuitionist approach to moral judgment. *Psychological review*, 108(4):814.
- Jonathan Haidt and Jesse Graham. 2007. When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research*, 20(1):98–116.
- Jonathan Haidt and Craig Joseph. 2004. Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus*, 133(4):55–66.

- Rujun Han, Qiang Ning, and Nanyun Peng. 2019. [Joint event and temporal relation extraction with shared representations and structured prediction](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 434–444, Hong Kong, China. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Joe Hoover, Kate Johnson, Reihane Boghrati, Jesse Graham, Morteza Dehghani, and M Brent Donnellan. 2018. Moral framing and charitable donation: Integrating exploratory social media analyses and confirmatory experimentation. *Collabra: Psychology*, 4(1).
- Joe Hoover, Gwenyth Portillo-Wightman, Leigh Yeh, Shreya Havaldar, Aida Mostafazadeh Davani, Ying Lin, Brendan Kennedy, Mohammad Atari, Zahra Kamel, Madelyn Mendlen, et al. 2020. Moral foundations twitter corpus: A collection of 35k tweets annotated for moral sentiment. *Social Psychological and Personality Science*, 11(8):1057–1071.
- Kristen Johnson and Dan Goldwasser. 2016. Identifying stance by analyzing political discourse on twitter. In *Proceedings of the First Workshop on NLP and Computational Social Science*, pages 66–75.
- Kristen Johnson and Dan Goldwasser. 2018. [Classification of moral foundations in microblog political discourse](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 720–730, Melbourne, Australia. Association for Computational Linguistics.
- Kristen Johnson and Dan Goldwasser. 2019. Modeling behavioral aspects of social media discourse for moral classification. In *Proceedings of the Third Workshop on Natural Language Processing and Computational Social Science*, pages 100–109.
- Ying Lin, Joe Hoover, Gwenyth Portillo-Wightman, Christina Park, Morteza Dehghani, and Heng Ji. 2018. Acquiring background knowledge to improve moral value prediction. In *2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 552–559. IEEE.
- Jiangming Liu, Shay B. Cohen, and Mirella Lapata. 2019. [Discourse representation parsing for sentences and documents](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6248–6262, Florence, Italy. Association for Computational Linguistics.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. A dataset for detecting stance in tweets. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 3945–3952.
- Marlon Mooijman, Joe Hoover, Ying Lin, Heng Ji, and Morteza Dehghani. 2018. Moralization in social networks and the emergence of violence during protests. *Nature human behaviour*, 2(6):389–396.
- Vlad Niculae, Joonsuk Park, and Claire Cardie. 2017. [Argument mining with structured SVMs and RNNs](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 985–995, Vancouver, Canada. Association for Computational Linguistics.
- Maria Leonor Pacheco and Dan Goldwasser. 2021. [Modeling Content and Context with Deep Relational Learning](#). *Transactions of the Association for Computational Linguistics*, 9:100–119.
- Chan Young Park, Xinru Yan, Anjalie Field, and Yulia Tsvetkov. 2020. Multilingual contextual affective analysis of lgbt people portrayals in wikipedia. *arXiv preprint arXiv:2010.10820*.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Daniel Preotjiuc-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.
- Shamik Roy and Dan Goldwasser. 2020. Weakly supervised learning of nuanced frames for analyzing polarization in news media. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7698–7716.
- Yanchuan Sim, Brice DL Acree, Justin H Gross, and Noah A Smith. 2013. Measuring ideological proportions in political speeches. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 91–101.
- Manuel Widmoser, Maria Leonor Pacheco, Jean Honorio, and Dan Goldwasser. 2021. [Randomized deep structured prediction for discourse-level processing](#). *Computing Research Repository*, arxiv:2101.10435.
- Jing Yi Xie, Renato Ferreira Pinto Junior, Graeme Hirst, and Yang Xu. 2019. Text-based inference of moral sentiment change. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4646–4655.
- Jerrold H Zar. 2005. Spearman rank correlation. *Encyclopedia of Biostatistics*, 7.

A Topic Indicator Lexicon

A.1 Topic Indicators for Gun Control

‘reduce gun’, ‘orlando shooting’, ‘terrorism watch’, ‘keep gun’, ‘terrorist watch’, ‘orlandounited’, ‘violence nobillnobreak’, ‘noflynobuy loophole’, ‘disarmhate’, ‘shooting’, ‘firearm’, ‘end gun’, ‘mas shooting’, ‘gun violence’, ‘sanbernadino’, ‘keeping gun’, ‘watch list’, ‘gun reform’, ‘hate crime’, ‘nobillnobreak’, ‘charleston9’, ‘gun safety’, ‘prevention legislation’, ‘gun owner’, ‘reducing gun’, ‘orlando terrorist’, ‘address gun’, ‘2nd amendment’, ‘gun show’, ‘tragic shooting’, ‘gun law’, ‘notonemore’, ‘ending gun’, ‘nomoresilence’, ‘closing terror’, ‘buy gun’, ‘nra’, ‘massacre’, ‘amendment right’, ‘reckles gun’, ‘endgunviolence’, ‘orlando terror’, ‘stopgunviolence’, ‘prevent gun’, ‘buying gun’, ‘gun loophole’, ‘gun legislation’, ‘massacred’, ‘sensible gun’, ‘sense gun’, ‘gun control’, ‘gun’, ‘terror watch’, ‘noflynobuy’, ‘standwithorlando’, ‘2a’, ‘charleston’, ‘gunviolence’, ‘background check’, ‘commonsense gun’, ‘guncontrol’

A.2 Topic Indicators for Immigration

‘fight for family’, ‘illegal immigrant’, ‘immigrant’, ‘granting amnesty’, ‘migration’, ‘asylum’, ‘dreamer’, ‘deportation’, ‘immigration action’, ‘homeland security’, ‘daca’, ‘fightforfamily’, ‘detain’, ‘borderwall’, ‘immigrationaction’, ‘border protection’, ‘daca work’, ‘sanctuarycity’, ‘sanctuary city’, ‘immigration detention’, ‘immigration system’, ‘immigration policy’, ‘illegal immigration’, ‘immigration’, ‘dacawork’, ‘detention’, ‘immigration reform’, ‘dhsgov’, ‘immigration law’, ‘executive amnesty’, ‘deport’, ‘dapa’, ‘immigration executive’, ‘refugee’, ‘border security’, ‘border wall’, ‘border sec’, ‘cir’, ‘comprehensive immigration’, ‘detained’, ‘detainee’, ‘amnesty’, ‘border protection’, ‘grant amnesty’, ‘deportee’, ‘immigr’

B Numeric Data of the Figure 2

The numeric values of each point in Figure 2 are as follows with standard deviations in brackets.

- Points fitting the red line in Figure 2(a): 1.0 (0), 0.889 (0.02), 0.880 (0.04), 0.897 (0.05), 0.864 (0.05)
- Points fitting the blue line in Figure 2(a): 0.864 (0.05), 0.913 (0.05), 0.960 (0.02), 0.954 (0.03), 1.0 (0)

- Points fitting the red line in Figure 2(b): 1.0 (0), 0.849 (0.02), 0.887 (0.03), 0.746 (0.03), 0.727 (0.04)
- Points fitting the blue line in Figure 2(b): 0.727 (0.04), 0.571 (0.04), 0.623 (0.03), 0.960 (0.01), 1.0 (0)

C Distribution of Most Polar Moral Foundation Usage over Grades

The distributions for the topics *Gun Control* and *Immigration* can be found in Figure 5 and Figure 6, respectively.

D Entities

D.1 Entities related to Gun Control

‘amendment’, ‘assault weapon ban’, ‘gun safety legislation’, ‘mexico’, ‘innocent life’, ‘gun sale’, ‘law enforcement’, ‘mass shooting’, ‘senseless gun violence’, ‘house judiciary’, ‘march life’, ‘young people’, ‘common sense gun reform’, ‘gun violence prevention’, ‘house gop’, ‘honor action’, ‘bump stock’, ‘wear orange’, ‘gun violence’, ‘assault weapon’, ‘republican’, ‘parkland’, ‘address gun violence’, ‘gun safety’, ‘gabby gifford’, ‘gun owner’, ‘las vegas’, ‘gun law’, ‘senate gop’, ‘mom demand’, ‘black’, ‘gun reform’, ‘tragic shooting’, ‘texas’, ‘dem’, ‘gun violence epidemic’, ‘congress’, ‘nra’, ‘police officer’, ‘town hall’, ‘virginia’, ‘bipartisan bill’, ‘pulse’, ‘universal background check’, ‘bipartisan background check’, ‘america’, ‘orlando’, ‘shannon r watt’, ‘end gun violence’, ‘school shooting’, ‘gun control’, ‘violence’, ‘american people’, ‘gun’, ‘community safe’, ‘el paso’, ‘high school’, ‘medicare’, ‘sandy hook’, ‘charleston’, ‘health care’, ‘gun lobby’, ‘background check’, ‘house democrat’

D.2 Entities related to Immigration

‘white house’, ‘hhs gov’, ‘republican’, ‘house judiciary’, ‘family’, ‘mexico’, ‘wall’, ‘refugee’, ‘supreme court’, ‘immigrant’, ‘protect dream’, ‘immigrant community’, ‘border patrol’, ‘dream act’, ‘protect dreamer’, ‘build wall’, ‘senate’, ‘american value’, ‘fema’, ‘human right’, ‘dreamer’, ‘save tps’, ‘asylum seeker’, ‘usc’, ‘illegal alien’, ‘hispanic caucus’, ‘immigration status’, ‘migrant child’, ‘ice’, ‘family separation’, ‘trump shutdown’, ‘detention facility’, ‘american citizen’, ‘homeland’, ‘real donald trump’, ‘ice gov’, ‘comprehensive immigration reform’, ‘dhs’, ‘illegal immigrant’, ‘defend daca’,




Figure 5: Moral Foundation distributions over NRA grades on Gun Control.

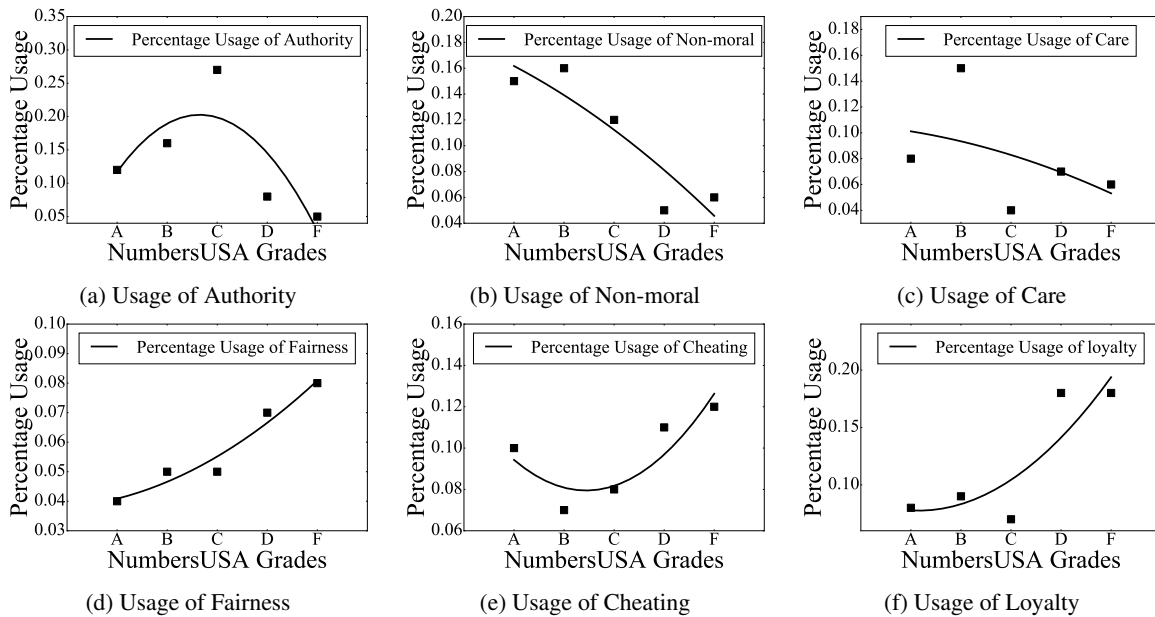


Figure 6: Moral Foundation distribution over NumbersUSA grades on Immigration.

‘family belong together’, ‘legal immigration’, ‘scotus’, ‘congress’, ‘daca’, ‘circuit court’, ‘government shutdown’, ‘muslim’, ‘dhs gov’, ‘immigration’, ‘national emergency’, ‘immigration system’, ‘immigration reform’, ‘border security’, ‘immigration law’, ‘immigrant family’, ‘anti immigrant agenda’, ‘house floor’, ‘america’, ‘c bp’, ‘sanctuary city’, ‘latino’, ‘humanitarian crisis’, ‘national security’, ‘dream promise’, ‘citizenship question’, ‘immigration policy’, ‘american people’, ‘border wall’, ‘detention center’, ‘dream promise act’, ‘southern border’, ‘immigrant child’, ‘medicare’, ‘keep fam-

ily together’, ‘illegal immigration’, ‘dream’, ‘circuit judge’, ‘young people’