

Ideological Phrase Indicators for Classification of Political Discourse Framing on Twitter

Kristen Johnson, I-Ta Lee, Dan Goldwasser

Department of Computer Science

Purdue University, West Lafayette, IN 47907

{john1187, lee2226, dgoldwas}@purdue.edu

Abstract

Politicians carefully word their statements in order to influence how others view an issue, a political strategy called framing. Simultaneously, these frames may also reveal the beliefs or positions on an issue of the politician. Simple language features such as unigrams, bigrams, and trigrams are important indicators for identifying the general frame of a text, for both longer congressional speeches and shorter tweets of politicians. However, tweets may contain multiple unigrams across different frames which limits the effectiveness of this approach. In this paper, we present a joint model which uses both linguistic features of tweets and ideological phrase indicators extracted from a state-of-the-art embedding-based model to predict the general frame of political tweets.

1 Introduction

Social media platforms have played an increasingly important role in U.S. presidential elections, beginning in 2008. Among these, microblogs such as Twitter have a special role, as they allow politicians to react quickly to events as they unfold and to shape the discussion of current political issues according to their views.

Framing is an important tool used by politicians to bias the discussion towards their stance. Framing contextualizes the discussion by emphasizing specific aspects of the issue, which creates an association between the issue and a specific frame of reference. Research on issue framing in political discourse is rooted in social science research (Entman, 1993; Chong and Druckman, 2007) and recently has attracted growing interest in the natural language processing community (Tsur et al.,

2015; Card et al., 2015; Baumer et al., 2015) as a way to automatically analyze political discourse in congressional speeches and political news articles. Contrary to these sources, Twitter requires politicians to compress their ideas and reactions into 140 character long tweets. As a result, politicians have to cleverly choose how to frame controversial issues, as well as react to events and each other (Mejova et al., 2013; Tumasjan et al., 2010).

Framing decisions can be used to build support for political stances and they often reflect ideological differences between politicians. For example, in debates concerning the issue of abortion, the stance opposing abortion is framed as “pro-life”, which reflects a moral or religious-based ideology. Correctly identifying how issues are framed can help reveal the ideological base of the speaker. However, in many cases framing abstracts this information and groups content reflecting differing ideologies together under the same frame. As a concrete example consider the following tweets:

1. *POTUS exec. order on guns is a gross overreach of power that tramples on the rights of law abiding Americans and our Constitution*
2. *With this ruling #SCOTUS has upheld a critical freedom for women to make their own decisions about their bodies*

In both tweets, the same frame (Legality, Constitutionality, & Jurisdiction) is used to discuss two different issues: guns and abortion, respectively. Despite the use of a similar frame, the two tweets reflect opposing ideologies.

A straight-forward approach for identifying these differences would be to refine the issue-independent general frames into more specific categories. However, this would limit their generalization and considerably increase the difficulty of analysis, both for human annotators and for automated techniques. Instead, we suggest to aug-

ment the frame analysis with additional information. Our modeling approach is based on the observation that politicians often use *slogans* in both their tweets and speeches. These are key phrases used to *indirectly* indicate the political figures' core beliefs and ideological stances. Identification of these phrases automatically decomposes the frames into more specific categories.

Consider the two tweets in the example above. In the first tweet, several phrases indicate the frame: "exec. order", "overreach of power", "rights of law abiding Americans", "our constitution". In the second tweet, the relevant phrases are "this ruling" and "upheld a critical freedom". All of these phrases indicate that the same frame is being used in both tweets. However, analyzing the specific terminology in each case and the context in which it appears helps capture the ideological similarities and differences. For example, in the context of gun-rights debates, phrases highlighting "law and order" and references to the constitution tend to reflect a conservative ideology, while phrases highlighting upholding of freedoms in the abortion debate tend to reflect a liberal ideology.

Given the rapidly changing nature of trending issues and political discourse on Twitter, our key technical challenge is to relay these ideological dimensions to an automated model, such that it will be able to easily adapt to new issues and language. Our model consists of two components combined together: frame identification and ideological-indicators identification. For the first piece we use a structured probabilistic model to capture general framing dimensions by combining content and political context analysis. For the second task, we employ a state-of-the-art textual similarity model which captures and generalizes over lexical indicators of key phrases that identify the politicians' ideology. More details of both components are described in Section 4.

In this paper we take a first step towards connecting these two dimensions of analysis: issue framing and ideology identification. We lay the foundation for more advanced research by identifying this connection, analyzing tweets authored by U.S congressional representatives, and extracting ideological phrase indicators. We build and analyze a joint model which combines the two dimensions. Our experiments in Section 5 quantitatively compare the differences in frame prediction performance when using ideological phrase indi-

cators. We also include a qualitative analysis in Section 6 of several examples in which ideological phrase indicators can help differentiate between tweets with similar frame predictions that reflect different ideologies.

2 Related Work

Previous computational works which analyze political discourse focus on opinion mining and stance prediction from forums and tweets (Sridhar et al., 2015; Hasan and Ng, 2014; Abu-Jbara et al., 2013; Walker et al., 2012; Abbott et al., 2011; Somasundaran and Wiebe, 2010, 2009; Johnson and Goldwasser, 2016; Ebrahimi et al., 2016). A variety of social media based predictions have been studied including: prediction of political affiliation and other demographics of Twitter users (Volkova et al., 2015, 2014; Yano et al., 2013; Conover et al., 2011), profile (Li et al., 2014b) and life event extraction (Li et al., 2014a), conversation modeling (Ritter et al., 2010), methods for handling unique microblog language (Eisenstein, 2013), and the modeling of social interactions and group structure in predictions (Sridhar et al., 2015; Abu-Jbara et al., 2013; West et al., 2014; Huang et al., 2012). Works which focus on inferring signed social networks (West et al., 2014) and collective classification using PSL (Bach et al., 2015) are similar to the modeling approach of Johnson et al. (2017b), which we extend in this paper.

Several previous works have explored framing in public statements, congressional speeches, and news articles (Fulgoni et al., 2016; Tsur et al., 2015; Card et al., 2015; Baumer et al., 2015). Framing is further related to works which analyze biased language (Recasens et al., 2013; Choi et al., 2012; Greene and Resnik, 2009) and subjectivity (Wiebe et al., 2004). Important to the language analysis of our work, Tan et al. (2014) have shown how wording choices can affect message propagation on Twitter. The study of political sentiment analysis (Pla and Hurtado, 2014; Bakliwal et al., 2013), ideology measurement and prediction (Iyyer et al., 2014; Bamman and Smith, 2015; Sim et al., 2013; Djemili et al., 2014), policies (Nguyen et al., 2015), voting patterns (Gerrish and Blei, 2012), and polls based on Twitter political sentiment (Birmingham and Smeaton, 2011; O'Connor et al., 2010; Tumasjan et al., 2010) are also related to the study of framing on Twitter.

FRAME NUMBER, FRAME, AND BRIEF DESCRIPTION
1. Economic: Economic effects of a policy
2. Capacity & Resources: Resources lack or availability
3. Morality & Ethics: Religious doctrine, righteousness, sense of responsibility
4. Fairness & Equality: Distribution of laws, punishments, resources, etc. among groups
5. Legality, Constitutionality, & Jurisdiction: Court cases and restriction and expressions of rights
6. Crime & Punishment: Crimes and consequences
7. Security & Defense: Preemptive actions to protect against threats
8. Health & Safety: Health care access and effectiveness
9. Quality of Life: Aspects of individual/community life
10. Cultural Identity: Trends, customs, and norms
11. Public Sentiment: Opinions and polling
12. Political Factors & Implications: Stances, filibusters, lobbying, references to political entities
13. Policy Description, Prescription, & Evaluation: Effectiveness of policies
14. External Regulation and Reputation: Interstate and international relationships
15. Factual: Expresses a fact, with no political spin
16. (Self) Promotion: Promotes author or another person
17. Personal Sympathy & Support: Expresses emotional response, including sympathy and solidarity

Table 1: General Frames and Their Descriptions. Detailed descriptions of the frames can be found in [Boydston et al. \(2014\)](#).

Political and social science works have studied the role of Twitter and framing in molding public opinion of events and issues ([Burch et al., 2015](#); [Harlow and Johnson, 2011](#); [Meraz and Papacharissi, 2013](#); [Jang and Hart, 2015](#)), as well as sentiment analysis and network agenda modeling of the 2012 U.S. presidential election ([Groshek and Al-Rawi, 2013](#)). [Boydston et al. \(2014\)](#) composed a Policy Frames Codebook for use in labeling general, issue-independent frames of longer texts. These frames were extended for Twitter and studied in a computational setting by [Johnson et al. \(2017b,a\)](#). Our approach builds upon these findings by identifying phrases which are relevant for determining ideology and increasing prediction accuracy of frames.

3 Data and Problem Setting

Dataset: In this work, we use the Congressional Tweets Dataset of [Johnson et al. \(2017b,a\)](#) which consists of the tweets of members of the 114th U.S. Congress. These tweets discuss six current political issues: (1) abortion, (2) the Affordable Care Act (i.e., the ACA or Obamacare), (3) gun ownership, (4) immigration, (5) terrorism, and (6) the LGBTQ community. The dataset provides a labeled portion of 2,050 tweets, which are labeled

using 17 possible frames. A brief description of each frame is shown in Table 1.

Frame Overlap: [Johnson et al. \(2017b,a\)](#) found that for most tweets, one or two frames were used. Additionally, in many cases, tweets authored by Republican and Democratic politicians use similar frames, both when discussing similar and different issues. For example, consider the following two tweets concerning the shooting of the Emanuel African Methodist Episcopal Church in 2015.

1. *Our thoughts and prayers must be with 9 innocent men and women murdered in Charleston, SC. Every effort must be made to capture the killer. RIP*
2. *My thoughts are with those impacted by the #CharlestonShooting. I pray that the perpetrator is brought to justice soon.*

Both tweets frame the shooting using two frames: Frame 6 (Crime & Punishment) and Frame 17 (Personal Sympathy & Support). In Tweet (1) the politician states that the killer must be captured. Similarly, in Tweet (2) the politician hopes for the perpetrator of the crime to be brought to justice. These phrases indicate that Frame 6 is being used. Additionally, in both tweets the politicians express that their thoughts are with those affected by the crime, indicating the use of Frame 17. Despite the use of the same frames by both tweets, there are very subtle differences between the two tweets, indicated by the specific phrase choices. For example, in Tweet (1) the politician uses the phrase “men and women murdered” to specifically reference the victims, while in Tweet (2) the politician uses “those impacted”, a more inclusive definition.

Phrase Identification: Using the labeled tweets of the dataset, we extracted lists of short phrases which frequently appear in each frame, for all frames.¹ All of these phrases can be further grouped into a more general phrase, which we term an *ideological phrase indicator*. For example, sub-phrases such as *rates will increase*, *increasing the rates this year*, and *premiums skyrocket* can be grouped into the more general ideological phrase indicator *Increase* of Frame 1 (Economic). From our observations, Democrats tend to

¹Phrases are currently extracted manually by matching them to the guidelines of [Boydston et al. \(2014\)](#). In future work, we aim to automate the phrase extraction.

Frame	General Ideological Phrase Indicators
Economic	Republican: Increase, Losses, Taxes, Job Effects Democrat: Deficit, Savings, Economy, Costs to Taxpayers
Capacity & Resources	Republican: Sources of Money, Defunding Democrat: Purchases, Taking Money Both Parties: Funding
Morality & Ethics	Republican: Morality Democrat: Sense of Obligation, Negative Descriptors Both Parties: Religion
Fairness & Equality	Republican: Race, Ethnicity Democrat: Women’s Rights, LGBT Rights, Discrimination, Civil Rights, Demands for Equality
Legality, Constitutionality, & Jurisdiction	Republican: Branches of Government Democrat: Items Being Voted On, SCOTUS Cases Both Parties: Laws, Rights
Crime & Punishment	Both Parties: Crimes
Security & Defense	Republican: Defense, Specific Threats Democrat: Ensure Safety, Preventive Measures Both Parties: Terrorism, Protection
Health & Safety	Republican: Health Care Aspects, Threats to Safety, Health Care Effectiveness Democrat: Health Insurance Access, Safety, Choices Both Parties: Health Care Access
Quality of Life	Republican: General Quality of Life Democrat: Affects Families, Affects Women’s Lives, Affects Everyone
Cultural Identity	Republican: Group Stereotypes Democrat: American, Immigrants Both Parties: Values
Public Sentiment	Both Parties: Americans Want, Polls
Political Factors & Implications	Both Parties: Republicans, Democrats, Congress, SCOTUS, POTUS
Policy Description, Prescription, & Evaluation	Republican: Votes on Bill Policies Democrat: Gun Policies, LGBT Policies, Immigration Policies Both Parties: ACA Policies, General Policies, Terrorism Policies
External Regulation & Reputation	Both Parties: National, International
Factual	Both Parties: Numerical Facts
(Self) Promotion	Both Parties: Media, References Self, References Others
Personal Sympathy & Support	Both Parties: Solidarity, Sympathy, Emotion

Table 2: Ideological Phrase Indicators for Each Frame. Frames are listed in the left column. General ideological phrase indicators used by each party, as well as by both parties, are listed in the right column.

use more phrase indicators (with more sub-phrases each) than Republicans for each frame. Finally, while the general phrase indicator name may be similar for both parties, the sub-phrases that are grouped under the general phrase may overlap, but are often different. For example, Frame 12 (Political Factors & Implications) has the general phrase indicator *Refers to POTUS* for both parties. However, the sub-phrases under this general phrase can differ across the parties, e.g. Republicans use phrases like “Obama admin” or “commander in chief”, while Democrats use phrases like “the administration”, “the president”, or “thank you POTUS”. Sub-phrases can also be similar across parties, e.g., both parties use “President Obama” in Frame 12. The general ideological phrase indica-

tors for each frame are listed in Table 2.²

4 PSL Models of Language on Twitter

Weakly Supervised Models with PSL: In order to model the dependencies between politicians and the language of their tweets, we design models with PSL, a declarative modeling language (Bach et al., 2015). PSL allows the user to specify first-order logic rules using domain knowledge. Weights for these rules are learned in either a supervised or unsupervised fashion and each weight indicates the importance of its associated rule. These rules are compiled into a hinge-loss Markov random field which defines a probability distribution over continuous value assignments to random variables of the model. For more details

²Complete lists of sub-phrases are omitted due to space.

PSL MODEL RULES
$\text{UNIGRAM}_F(T, U) \wedge \text{SIMPHRASE}(T, P_F) \rightarrow \text{FRAME}(T, F)$
$\text{UNIGRAM}_F(T, U) \wedge \text{PARTY}(T, P) \wedge \text{SIMPHRASE}(T, P_F) \rightarrow \text{FRAME}(T, F)$
$\text{UNIGRAM}_F(T, U) \wedge \text{SIMUNIGRAM}(T, F) \wedge \text{SIMPHRASE}(T, P_F) \rightarrow \text{FRAME}(T, F)$
$\text{UNIGRAM}_F(T, U) \wedge \text{PARTY}(T, P) \wedge \text{BIGRAM}_P(T, B) \wedge \text{SIMPHRASE}(T, P_F) \rightarrow \text{FRAME}(T, F)$
$\text{UNIGRAM}_F(T, U) \wedge \text{PARTY}(T, P) \wedge \text{TRIGRAM}_P(T, TR) \wedge \text{SIMPHRASE}(T, P_F) \rightarrow \text{FRAME}(T, F)$

Table 3: Examples of PSL Model Rules. Predicates composed into rules are on the left hand side and the target predicate (prediction goal) is on the right hand side.

on PSL we refer the reader to [Bach et al. \(2015\)](#).

To evaluate if modeling ideological phrase indicators can increase the F_1 score of frame prediction, we use the most indicative features for predicting a tweet’s frame (as determined by [Johnson et al. \(2017b\)](#)): unigrams, word similarity to unigrams, bigrams, and trigrams. In addition, we add tweet similarity to phrases ($\text{SIMPHRASE}(T, P_F)$ described below) as a feature. These features are extracted using weakly supervised models and represented as the following predicates in PSL notation: $\text{UNIGRAM}_F(T, U)$, $\text{SIMUNIGRAM}(T, F)$, $\text{BIGRAM}_P(T, B)$, $\text{TRIGRAM}_P(T, TR)$. Each predicate indicates that the tweet T has that unigram U , a word similar to that unigram, a bigram B , or a trigram TR , respectively. Finally, the party of the politician who authored the tweet ($\text{PARTY}(T, P)$) is also used. These predicates are combined into the probabilistic rules of the PSL model as shown in Table 3.

Incorporating Phrase Similarity: Due to the dynamic nature of language and trending political issues on Twitter, it is infeasible to construct a list of all possible phrases one can expect politicians to use when framing an issue. Therefore, we use the embedding-based model of [Lee et al. \(2017\)](#) to determine which tweets contain phrases that are similar to our initial list of phrases. For example, given the phrase *insurance rates will increase*, we want to find all tweets which contain similar phrases, e.g., *rising insurance premiums*.

The phrase similarity model was trained on the Paraphrase Database (PPDB) ([Ganitkevitch et al., 2013](#)) and incorporates a Convolutional Neural Network (CNN) to capture sentence structures. This model generates the embeddings of our phrases and computes the cosine similarities between phrases and tweets as the scores. The input tweets and phrases are represented as the average word embeddings in the input layer, which are then projected into a convolutional layer, a max-pooling layer, and finally two fully-connected lay-

ers. The embeddings are thus represented in the final layer. The learning objective of this model is:

$$\min_{W_c, W_w} \left(\sum_{\langle x_1, x_2 \rangle \in X} \max(0, \delta - \cos(g(x_1), g(x_2))) \right. \\ \left. + \cos(g(x_1), g(t_1)) \right) \\ + \max(0, \delta - \cos(g(x_1), g(x_2))) \\ \left. + \cos(g(x_2), g(t_2)) \right) \\ + \lambda_c \|W_c\|^2 + \lambda_w \|W_{init} - W_w\|^2,$$

where X is all the positive input pairs, δ is the margin, $g(\cdot)$ represents the network, λ_c and λ_w are the weights for L2-regularization, W_c is the network parameters, W_w is the word embeddings, W_{init} is the initial word embeddings, and t_1 and t_2 are negative examples that are randomly selected.

All tweet-phrase pairs with a cosine similarity over a given threshold are used as input to the PSL model via the predicate $\text{SIMPHRASE}(T, P_F)$, which indicates that tweet T contains a phrase that is similar to the phrases for a certain frame (P_F). Table 3 presents examples of the rules used in our modeling procedure.

5 Experiments

Analysis of Supervised Experiments: Since each tweet can be classified as having more than one frame, the prediction task becomes a multi-label classification task. Therefore, we use the standard measurements for precision and recall of a multilabel task. The F_1 score is the harmonic mean of these two measures. We conducted supervised experiments using five-fold cross validation with randomly chosen splits on the labeled portion of the dataset. Table 4 shows the results of our supervised experiments. The first column lists the frame number. The second column presents the results of the baseline model, which includes all of the rules listed in Table 3 *without* the $\text{SIMPHRASE}(T, P_F)$ predicate. The third

FRAME NO.	BASELINE	PHRASES
1	85.11	87.50
2	82.35	82.05
3	88.46	76.79
4	82.35	75.28
5	67.57	71.57
6	63.64	70.59
7	83.12	89.70
8	75.68	89.51
9	76.47	71.52
10	88.89	84.52
11	29.41	29.63
12	73.92	81.25
13	65.43	62.35
14	85.71	82.25
15	82.35	83.33
16	82.05	73.55
17	91.07	91.67
Weighted Avg.	75.95	76.27

Table 4: F₁ Scores of Supervised Experiments. The baseline column represents the results of the best model of Johnson et al. (2017b). The phrases column indicates the scores for the best model when combined with our proposed phrases. Items in bold are the highest score. The weighted average is the micro-weighted average of the F₁ scores.

column lists the results of our model which consists of the baseline model *with* the addition of the SIMPHRASE(T,P_F) predicate.

From these results we can see that the joint model that uses both language features (i.e., unigrams, bigrams, and trigrams) and phrase indicators (shown in Table 2) is able to improve performance in 9 out of the 17 frames. The most likely cause for the decrease in score for the other 8 frames is that it is possible that there are too many overlapping sub-phrases within the general phrases of these 8 frames. This would introduce extra noise into the probabilistic model and result in lower scores. The 9 frames which improve have either 1 or no overlapping sub-phrases across parties for each general phrase category. Further refinement of the sub-phrases is left for future work.

Ablation Case Study: To investigate the usefulness of ideological phrase indicators, we conducted an ablation study on the results of Frame 12. Frame 12 is used when a politician references other political entities (e.g., the House, Senate, former presidents, etc.) as well as political actions (e.g., filibusters or lobbying). For our dataset, we used the following general phrases for Frame 12 which include references to: Democrats, Republicans, the President (POTUS), the Supreme Court

(SCOTUS), and Congress. We ran our model through an ablation study, in which each pair of phrases is removed one at a time to study their overall effect on the final prediction. Table 5 presents the results of this experiment.

MODEL	F ₁ SCORE	CHANGE
All Phrases	81.25	—
Republicans	85.71	+ 4.46
Democrats	77.78	- 3.47
POTUS	83.33	+ 2.08
SCOTUS	85.71	+ 4.46
Congress	78.57	- 2.68

Table 5: F₁ Scores of Ablation Experiments. All Phrases represents our score for Frame 12 when using all possible phrases. The remaining rows indicate which general phrase indicators have been removed from the comprehensive model. Column 2 presents the F₁ score. Column 3 indicates the increase or decrease in score after the respective phrases are removed.

From these initial results, it appears that the way politicians refer to Democrats and Congress are the most important phrase indicators for predicting Frame 12. When these two phrase groups are removed, there is a large decrease in F₁ score. Additionally, removing references to the president has a slight increase, while removing references to Republicans and the Supreme Court has a larger increase. Therefore, references to Republicans and the Supreme Court are likely to be the least useful for predicting this frame. We leave finding the best combinations of phrases for each frame as future work, as described in Section 7.

6 Qualitative Analysis

The supervised experiments of the previous section allow us to analyze the effects of phrases as features for frame prediction. In this section, we explore the predictions of the phrase-based model to locate framing trends of a real world event. We first learned the weights of each model using the labeled data and then performed MPE inference on the *unlabeled* tweets to obtain their predicted frames. We used these predictions to analyze the political discourse on Twitter by focusing on tweets concerning the shooting of the Pulse Nightclub in Orlando, Florida (June 12, 2016). Table 6 presents the frame predictions and example tweets for this event.

Frame 17 reflects politicians tweeting that their

DATE	POLITICIAN	POLITICAL PARTY	TWEET	PREDICTED FRAME(S)
6/12/2016	Alex Mooney	Republican	My thoughts and prayers are with the people of #Orlando, the victims, and their families.	17
6/12/2016	Brad Ashford	Democrat	As authorities investigate the Orlando shooting, we must pray for the victims and act swiftly to keep these tragedies out of our communities.	9
6/12/2016	Lisa Murkowski	Republican	What happened in Orlando was an absolute tragic act of terrorism spawned by an ideology of hate being pushed by ISIS.	3
6/12/2016	Bob Goodlatte	Republican	The attack in #Orlando was an act of pure evil. My prayers are w/ the families of victims and the injured. We will continue seeking answers.	3, 17
6/12/2016	David Cicilline	Democrat	Voters should absolutely hold us accountable for what we're doing or not doing to address gun violence.	3
6/12/2016	Yvette Clark	Democrat	I am deeply saddened by the act of hate and terror enacted on the lives of Orlando's LGBT Community and I #StandWithOrlando	3, 17
6/15/2016	Jeanne Shaheen	Democrat	Joining @ChrisMurphyCT on the Senate floor to say #Enough and call for reforms 2 prevent gun violence.	7, 12
6/15/2016	Mark Kirk	Republican	Americans need to know Washington is listening - We must keep guns out of the hands of suspected terrorists	7
6/15/2016	Kirsten Gillibrand	Democrat	As we mourn victims of yet another tragedy, time to finally act on commonsense gun safety reforms supported by the American people.	11, 12

Table 6: Example Tweets Associated With the Orlando Pulse Nightclub Shooting on June 12, 2016.

“thoughts and prayers” are with the community, as seen in the first line of Table 6. Offers of prayers and sympathy are used by both parties as the initial response the day this (and most other) shootings occur. This can be considered both as a reflection of the politicians’ immediate emotional reaction to the shooting but also to support other agendas, as Frame 17 also appears in tweets that use other frames, specifically Frames 9 and 3. Interestingly, Republicans and Democrats use these frames in nuanced ways to promote different agendas, which are identifiable by the presence (or lack thereof) of different key phrases.

Republicans used Frame 3, often in combination with Frame 17, to discuss the shooting as an act of evil or terrorism as well as to suggest links between the shooter and ISIS (examples of these tweets are shown in rows three and four of Table 6). Democrats, however, used Frame 3 to express a sense of responsibility on their part to take actions to prevent gun violence (e.g., row five of Table 6) or refer to the shooting as a hate crime or act of terror (e.g., row six of Table 6). All of these examples are expressed with Frame 3, however, the different phrases indicate differing underlying ideologies. For example, referring to the shooting as an “act of evil” indicates a religious-based ideology, which also limits possible ways to combat the problem. However, by associating the

cause with hatred or terror instead, there is a subtle implication that measures can be taken to prevent future violence with similar causes. Democrats go one step further by using this frame to transition into calls for increased gun legislation, which would be a concrete step towards preventing future shootings.

On June 15th, three days after the shooting, Democrats held a filibuster to push for a vote on gun control. The top frame that day for both parties is Frame 7 (Security & Defense), however different phrases represent different ideologies in this example as well. Democrats frame the need for gun control laws as a preemptive measure that will prevent gun violence (e.g., row seven of Table 6). Republicans use Frame 7 to discuss the need to prevent threats posed by ISIS (possibly due to the shooter’s association with ISIS) as shown in row eight of Table 6. Additionally, some Republicans promote bipartisan efforts to stop the sale of guns to known terrorists (row eight). While all examples use Frame 7 to support gun control, this support is limited depending on party and identifiable by different key phrases, e.g. the general goal of “reforms 2 prevent gun violence” versus the specific target to “keep guns out of the hands of suspected terrorists”.

Lastly, the impacts of the shooting on the quality of life of the community (or nation as a whole)

are discussed in tweets having Frame 9. For example, row two of Table 6 shows a Democrat’s tweet calling for action to keep gun violence tragedies from affecting communities. For this event, Republicans are more likely to refer to the “Orlando community” while Democrats are more likely to reference the “LGBT community”, indicating that national versus specific-group phrases are useful in identifying Frame 9.

7 Future Work

Currently, this work requires human knowledge and engineering to compile the sub-phrases by party. Additionally, for computational simplicity all phrases are currently added to the baseline model for evaluation. Since frames can overlap and politicians can use the talking points of other parties, we hypothesize that frame prediction can be further improved by automatically testing all possible phrases with the baseline model.

For future work, we are building an automatic search over all possible phrase indicators, designed to choose the most indicative phrases for predicting each frame. We hope this tool will be useful for scientists from other fields, allowing them to compile their expert knowledge of a domain into many rules, which can then be analyzed to indicate the most useful features for further study of a subject.

8 Conclusion

In this paper we present an analysis of the usefulness of ideological phrases as a feature for predicting the frame of a political tweet. By compiling a list of common phrases and computing their similarity to tweets, we are able to increase the F_1 scores for half of the frames over a simpler language based model. We provide an analysis of our joint model in a supervised setting and show interesting real world examples. Finally, we propose the automation of phrase searching as a future work to improve the usefulness of this technique in other scientific communities.

Acknowledgments

We thank the anonymous reviewers for their thoughtful comments and suggestions.

References

- Rob Abbott, Marilyn Walker, Pranav Anand, Jean E. Fox Tree, Robeson Bowmani, and Joseph King. 2011. How can you say such things?!?: Recognizing disagreement in informal political argument. In *Proc. of the Workshop on Language in Social Media*.
- Amjad Abu-Jbara, Ben King, Mona Diab, and Dragomir Radev. 2013. Identifying opinion subgroups in arabic online discussions. In *Proc. of ACL*.
- Stephen H Bach, Matthias Broecheler, Bert Huang, and Lise Getoor. 2015. Hinge-loss markov random fields and probabilistic soft logic. *arXiv preprint arXiv:1505.04406*.
- Akshat Bakliwal, Jennifer Foster, Jennifer van der Puil, Ron O’Brien, Lamia Tounsi, and Mark Hughes. 2013. Sentiment analysis of political tweets: Towards an accurate classifier. In *Proc. of ACL*.
- David Bamman and Noah A Smith. 2015. Open extraction of fine-grained political statements. In *Proc. of EMNLP*.
- Eric Baumer, Elisha Elovic, Ying Qin, Francesca Polletta, and Geri Gay. 2015. Testing and comparing computational approaches for identifying the language of framing in political news. In *Proc. of NAACL*.
- Adam Bermingham and Alan F Smeaton. 2011. On using twitter to monitor political sentiment and predict election results.
- 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.
- Lauren M. Burch, Evan L. Frederick, and Ann Pegoraro. 2015. Kissing in the carnage: An examination of framing on twitter during the vancouver riots. *Journal of Broadcasting & Electronic Media* 59(3):399–415.
- Dallas Card, Amber E. Boydston, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. The media frames corpus: Annotations of frames across issues. In *Proc. of ACL*.
- Eunsol Choi, Chenhao Tan, Lillian Lee, Cristian Danescu-Niculescu-Mizil, and Jennifer Spindel. 2012. Hedge detection as a lens on framing in the gmo debates: A position paper. In *Proc. of ACL Workshops*.
- Dennis Chong and James N Druckman. 2007. Framing theory. *Annu. Rev. Polit. Sci.* 10:103–126.
- Michael D Conover, Bruno Gonçalves, Jacob Ratkiewicz, Alessandro Flammini, and Filippo Menczer. 2011. Predicting the political alignment of twitter users. In *Proc. of PASSAT*.

- Sarah Djemili, Julien Longhi, Claudia Marinica, Dimitris Kotzinos, and Georges-Elia Sarfati. 2014. What does twitter have to say about ideology? In *NLP 4 CMC*.
- Javid Ebrahimi, Dejing Dou, and Daniel Lowd. 2016. Weakly supervised tweet stance classification by relational bootstrapping. In *Proc. of EMNLP*.
- Jacob Eisenstein. 2013. What to do about bad language on the internet. In *Proc. of NAACL*.
- Robert M Entman. 1993. Framing: Toward clarification of a fractured paradigm. *Journal of communication* 43(4):51–58.
- Dean Fulgoni, Jordan Carpenter, Lyle Ungar, and Daniel Preotiuc-Pietro. 2016. An empirical exploration of moral foundations theory in partisan news sources. In *Proc. of LREC*.
- Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2013. The paraphrase database. In *Proc. of NAACL-HLT*.
- Sean Gerrish and David M Blei. 2012. How they vote: Issue-adjusted models of legislative behavior. In *Advances in Neural Information Processing Systems*, pages 2753–2761.
- Stephan Greene and Philip Resnik. 2009. More than words: Syntactic packaging and implicit sentiment. In *Proc. of NAACL*.
- Jacob Groshek and Ahmed Al-Rawi. 2013. Public sentiment and critical framing in social media content during the 2012 u.s. presidential campaign. *Social Science Computer Review* 31(5):563–576.
- Summer Harlow and Thomas Johnson. 2011. The arab spring— overthrowing the protest paradigm? how the new york times, global voices and twitter covered the egyptian revolution. *International Journal of Communication* 5(0).
- Kazi Saidul Hasan and Vincent Ng. 2014. Why are you taking this stance? identifying and classifying reasons in ideological debates. In *Proc. of EMNLP*.
- Bert Huang, Stephen H. Bach, Eric Norris, Jay Pujara, and Lise Getoor. 2012. Social group modeling with probabilistic soft logic. In *NIPS Workshops*.
- Iyyer, Enns, Boyd-Graber, and Resnik. 2014. Political ideology detection using recursive neural networks. In *Proc. of ACL*.
- S. Mo Jang and P. Sol Hart. 2015. Polarized frames on "climate change" and "global warming" across countries and states: Evidence from twitter big data. *Global Environmental Change* 32:11–17.
- Kristen Johnson and Dan Goldwasser. 2016. All i know about politics is what i read in twitter: Weakly supervised models for extracting politicians' stances from twitter. In *Proc. of COLING*.
- Kristen Johnson, Di Jin, and Dan Goldwasser. 2017a. Leveraging behavioral and social information for weakly supervised collective classification of political discourse on twitter. In *Proc. of ACL*.
- Kristen Johnson, Di Jin, and Dan Goldwasser. 2017b. Modeling of political discourse framing on twitter. In *Proc. of ICWSM*.
- I-Ta Lee, Mahak Goindani, Chang Li, Di Jin, Kristen Johnson, Xiao Zhang, Maria Pacheco, and Dan Goldwasser. 2017. Purduenlp at semeval-2017 task 1: Predicting semantic textual similarity with paraphrase and event embeddings. In *Proc. of SemEval*.
- Jiwei Li, Alan Ritter, Claire Cardie, and Eduard H Hovy. 2014a. Major life event extraction from twitter based on congratulations/condolences speech acts. In *Proc. of EMNLP*.
- Jiwei Li, Alan Ritter, and Eduard H Hovy. 2014b. Weakly supervised user profile extraction from twitter. In *Proc. of ACL*.
- Mejova, Srinivasan, and Boynton. 2013. Gop primary season on twitter: popular political sentiment in social media. In *WSDM*.
- Sharon Meraz and Zizi Papacharissi. 2013. Networked gatekeeping and networked framing on #egypt. *The International Journal of Press/Politics* 18(2):138–166.
- Viet-An Nguyen, Jordan 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 *Proc. of ACL*.
- Brendan O'Connor, Ramnath Balasubramanyan, Bryan R Routledge, and Noah A Smith. 2010. From tweets to polls: Linking text sentiment to public opinion time series. In *Proc. of ICWSM*.
- Ferran Pla and Lluís F Hurtado. 2014. Political tendency identification in twitter using sentiment analysis techniques. In *Proc. of COLING*.
- Marta Recasens, Cristian Danescu-Niculescu-Mizil, and Dan Jurafsky. 2013. Linguistic models for analyzing and detecting biased language. In *Proc. of ACL*.
- Alan Ritter, Colin Cherry, and Bill Dolan. 2010. Unsupervised modeling of twitter conversations. In *Proc. of NAACL*.
- Sim, Acree, Gross, and Smith. 2013. Measuring ideological proportions in political speeches. In *Proc. of EMNLP*.
- Swapna Somasundaran and Janyce Wiebe. 2009. Recognizing stances in online debates. In *Proc. of ACL*.
- Swapna Somasundaran and Janyce Wiebe. 2010. Recognizing stances in ideological on-line debates. In *Proc. of NAACL Workshops*.

- Dhanya Sridhar, James Foulds, Bert Huang, Lise Getoor, and Marilyn Walker. 2015. Joint models of disagreement and stance in online debate. In *Proc. of ACL*.
- Chenhao Tan, Lillian Lee, and Bo Pang. 2014. The effect of wording on message propagation: Topic- and author-controlled natural experiments on twitter. In *Proc. of ACL*.
- Oren Tsur, Dan Calacci, and David Lazer. 2015. A frame of mind: Using statistical models for detection of framing and agenda setting campaigns. In *Proc. of ACL*.
- Andranik Tumasjan, Timm Oliver Sprenger, Philipp G Sandner, and Isabell M Welpe. 2010. Predicting elections with twitter: What 140 characters reveal about political sentiment. In *Proc. of ICWSM*.
- Svitlana Volkova, Yoram Bachrach, Michael Armstrong, and Vijay Sharma. 2015. Inferring latent user properties from texts published in social media. In *Proc. of AAAI*.
- Svitlana Volkova, Glen Coppersmith, and Benjamin Van Durme. 2014. Inferring user political preferences from streaming communications. In *Proc. of ACL*.
- Marilyn A. Walker, Pranav Anand, Robert Abbott, and Ricky Grant. 2012. Stance classification using dialogic properties of persuasion. In *Proc. of NAACL*.
- Robert West, Hristo S Paskov, Jure Leskovec, and Christopher Potts. 2014. Exploiting social network structure for person-to-person sentiment analysis. *TACL*.
- Janyce Wiebe, Theresa Wilson, Rebecca Bruce, Matthew Bell, and Melanie Martin. 2004. Learning subjective language. *Computational linguistics*.
- Tae Yano, Dani Yogatama, and Noah A Smith. 2013. A penny for your tweets: Campaign contributions and capitol hill microblogs. In *Proc. of ICWSM*.