The Pragmatics behind Politics: Modelling Metaphor, Framing and Emotion in Political Discourse

Pere-Lluís Huguet Cabot[1], **Verna Dankers**[1,3], **David Abadi**[2], **Agneta Fischer**[2], **Ekaterina Shutova**[1]
[1] Institute for Logic, Language and Computation, University of Amsterdam
[2] Department of Psychology, University of Amsterdam
[3] Institute for Language, Cognition and Computation, University of Edinburgh
{perelluis1993, vernadankers}@gmail.com,
{d.r.abadi, A.H.Fischer, e.shutova}@uva.nl

Abstract

There has been an increased interest in modelling political discourse within the natural language processing (NLP) community, in tasks such as political bias and misinformation detection, among others. Metaphor-rich and emotion-eliciting communication strategies are ubiquitous in political rhetoric, according to social science research. Yet, none of the existing computational models of political discourse has incorporated these phenomena. In this paper, we present the first joint models of metaphor, emotion and political rhetoric, and demonstrate that they advance performance in three tasks: predicting political perspective of news articles, party affiliation of politicians and framing of policy issues.

1 Introduction

The role of metaphor and emotion in political discourse has been investigated in fields such as communication studies (Weeks, 2015; Mourão and Robertson, 2019), political science (Ferrari, 2007; Charteris-Black, 2009) and psychology (Edwards, 1999; Bougher, 2012). Political rhetoric may rely on metaphorical framing to shape public opinion (Lakoff, 1991; Musolff, 2004). Framing selectively emphasises certain aspects of an issue that promote a particular perspective (Entman, 1993). For instance, government spending on the wealthy can be portrayed as a *partnership* or *bailout*, spending on the middle class as simply *spending* or *stimulus to the economy* and spending on the poor as a *giveaway* or a *moral duty*, the former corresponding to the conservative and the latter to the liberal point of view (Stone, 1988). Metaphor is an apt framing device, with different metaphors used across communities with distinct political views (Kövecses, 2002; Lakoff and Wehling, 2012). At the same time, metaphorical language has been shown to express and elicit stronger emotion than literal language (Citron and Goldberg, 2014; Mohammad et al., 2016) and to provoke emotional responses in the context of political discourse covered by mainstream newspapers (Figar, 2014). For instance, the phrase "immigrants are *strangling* the welfare system" aims to promote fear of immigration. On the other hand, the experienced emotions may influence the effects of news framing on public opinions (Lecheler et al., 2015) and individual variations in emotion regulation styles can predict different political orientations and support for conservative policies (Lee Cunningham et al., 2013). Metaphor and emotion thus represent crucial tools in political communication.

At the same time, computational modelling of political discourse, and its specific aspects, such as political bias in news sources (Kiesel et al., 2019), framing of societal issues (Card et al., 2015), or prediction of political affiliation from text (Iyyer et al., 2014) have received a great deal of attention in the NLP community. Yet, none of this research has incorporated the notions of metaphor and emotion in modelling political rhetoric.

We present the first joint models of metaphor, emotion and political rhetoric, within a multi-task learning (MTL) framework. We make use of auxiliary learning, i.e. training a model in more than one task to improve the performance on a main task. We experiment with three tasks from the political realm, predicting (1) political perspective of a news article; (2) party affiliation of politicians from their social media posts; and (3) framing dimensions of policy issues. We use metaphor and emotion detection as auxiliary tasks, and investigate whether incorporating metaphor or emotion-related features enhances the models of political discourse. Our results show that incorporating metaphor or emotion significantly improves performance across all tasks, emphasising the prominent role they play in political rhetoric.

Findings of the Association for Computational Linguistics: EMNLP 2020, pages 4479–4488

2 Related work

Modelling political discourse encompasses a broad spectrum of tasks, including estimating policy positions from political texts (Thomas et al., 2006; Lowe et al., 2011), identifying features that differentiate political rhetoric of opposing parties (Monroe et al., 2008) or predicting political affiliation of Twitter users (Conover et al., 2011; Pennacchiotti and Popescu, 2011; Preoţiuc-Pietro et al., 2017). Deep neural networks have been widely used to model political perspective, bias or affiliation at document level: Iyyer et al. (2014) used a recurrent neural network (RNN) to predict political affiliation from US congressional speeches. Li and Goldwasser (2019) identified the political perspective of news articles using a hierarchical Long Short-Term Memory (LSTM) and modelled social media user data with Graph Convolutional Networks (GCN). Lastly, a recent shared task presented a multitude of deep learning methods to detect political bias in articles (Kiesel et al., 2019). Framing in political discourse has gained some attention recently. Ji and Smith (2017) trained Tree-RNNs to classify framing of policy issues in news articles. Johnson et al. (2017) used a weakly supervised approach to identify tweet-level frames and used Probabilistic Soft Logic on language and social-based features. Hartmann et al. (2019) classified frames at a sentence level using bidirectional LSTMs and GRUs and recently Khanehzar et al. (2019) compared a set of classifiers on frame identification in news.

Approaches predicting emotions for a given text typically adopt a categorical model of discrete, prototypical emotions, e.g. the six basic emotions of Ekman (1992). Early computational approaches employed vector space models (Danisman and Alpkocak, 2008) or shallow machine learning classifiers (Alm et al., 2005; Yang et al., 2007). Examples of deep neural methods are the recurrent model of Abdul-Mageed and Ungar (2017), who classified 24 fine-grained emotions, and the transformer-based SentiBERT architecture of Yin et al. (2020).

Computational research on metaphor has mainly focused on detecting metaphorical language in text. Early research performed supervised classification with hand-engineered lexical, syntactic and psycholinguistic features (Tsvetkov et al., 2014; Beigman Klebanov et al., 2016; Turney et al., 2011; Strzalkowski et al., 2013; Bulat et al., 2017). Alternative approaches perform metaphor detection from distributional properties of words (Shutova et al., 2010; Gutiérrez et al., 2016) or by training deep neural models (Rei et al., 2017; Gao et al., 2018). Dankers et al. (2019) developed a joint model of metaphor and emotion by fine-tuning BERT in an MTL setting.

3 Tasks and datasets

Political Perspective in News Media Political news can be biased towards the left or right side of the political spectrum. To model such biased perspectives computationally, we classify articles as *left*, *right* or *centre* using data from Li and Goldwasser (2019).[1] The articles are from the website AllSides[2] and are annotated with their source's bias. The training and test sets contain 2008 and 5761 articles, respectively. We use 30% of training data for validation.

Political Affiliation For this task, we use the dataset of Voigt et al. (2018),[3] which contains public Facebook posts from 412 US politicians. The training, validation and test sets contain 9792, 2356 and 2458 posts, respectively. The classes are balanced and each set does not share politicians with the other sets. The task is thus to predict *republican* or *democrat* for posts of unseen politicians.

Framing The Media Frames Corpus[4] (Card et al., 2015) contains news articles discussing five policy issues: tobacco, immigration, same-sex marriage, gun-control and death penalty. There are 15 possible framing dimensions, e.g. *economic*, *political* etc. (see Appendix A.3.1). We use the article-level annotation to predict the framing dimension. Of 23,580 articles, we use 15% as the test set, and 15% of the training data for validation.

Metaphor For metaphor detection we use the VU Amsterdam dataset (Steen et al., 2010), which is a subset of the British National Corpus (Leech, 1992). The dataset contains 9,017 sentences and binary labels (*literal* or *metaphorical*) per word. We use the data split of Gao et al. (2018), that includes 25% of the sentences in the test set.

Emotion For emotion classification, we use a dataset from SemEval-2018 Task 1 (Mohammad et al., 2018), in which tweets were labelled for eleven emotion classes or as neutral (see Ap-

[1] https://github.com/BillMcGrady/NewsBiasPrediction
[2] https://www.allsides.com/unbiased-balanced-news
[3] https://nlp.stanford.edu/robvoigt/rtgender/
[4] https://github.com/dallascard/media_frames_corpus

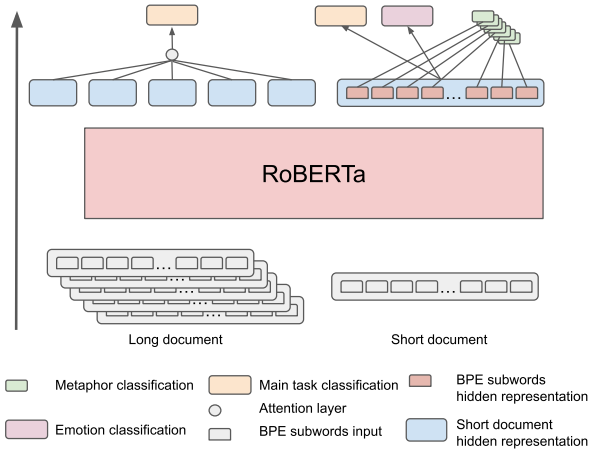

Figure 1: Schematics of the MTL model. The left side shows the path for longer documents from the *Political Perspective in News Media* dataset, while the right side is the path for the rest of datasets and the auxiliary tasks.

pendix A.3.2). We use the English portion of the dataset (10,097 tweets) and the shared task splits.

4 Methods

We employ the Robustly optimized BERT approach (RoBERTa-base) presented by Liu et al. (2019) and use the implementation by Wolf et al. (2019). RoBERTa contains twelve stacked transformer layers and assumes an input sequence to be tokenised into subword units called Byte-Pair Encodings (BPE). A special <s> token is inserted at the beginning of the input sequence, to compute a contextualised sequence representation.

Our tasks are defined at three levels of the linguistic hierarchy. The auxiliary tasks of metaphor detection and emotion prediction are defined at word and sentence level, respectively, while the main political tasks are defined at document level.

For word-level metaphor identification, the subword encodings from RoBERTa's last layer are processed by a linear classification layer. A word is considered metaphorical provided that any of its BPEs was labelled as metaphorical. We assume the BPE from inflexions unlikely to cause a word to be incorrectly labelled as metaphorical.

For the sentence-level emotion prediction task and the document-level tasks of political affiliation and framing, the <s> encoding serves as sequence representation and is fed to a linear classification layer. For political perspective in news articles, the average document length exceeds the maximum input size of RoBERTa. We, therefore, split its documents into sentences and collect them in a maximum of 5 subdocuments with up to 256 subwords. After applying RoBERTa to the subdocuments, their <s> encodings are fed to an attention layer yielding a document representation to be classified. A model schematic is shown in Figure 1.

All task models use the cross-entropy loss with a sigmoid activation function. For political perspective detection, the loss function includes class weights to account for class imbalance.

4.1 Multi-task learning

The MTL architecture uses hard parameter sharing for the first eleven transformer layers. The last layer of RoBERTa, the classification and attention layers are task-specific to allow for specialisation, similar to the approach of Dankers et al. (2019).

The main political tasks are paired with the metaphor and emotion tasks one by one. The task losses are weighted with α for the main task and $1-\alpha$ for the auxiliary task. We include an auxiliary warm-up period, during which $\alpha = 0.01$, for some tasks. This allows the model to initially learn the (lower-level) auxiliary task while focusing mostly on the main task afterwards. This approach is similar to Kiperwasser and Ballesteros (2018).

5 Experiments and results

5.1 Experimental setup

The models are trained with the *AdamW* optimiser, a learning rate of $1e-5$ and a batch size of 32. The learning rate is annealed through a cosine-based schedule and warm-up ratios of 0.2, 0.3 and 0.15 for political perspective in news media, the political affiliation and the framing tasks, respectively. Dropout is applied per layer with a probability of 0.3 for political affiliation and 0.1 otherwise.

The auxiliary warm-up period and α values are estimated per main task, for metaphor (α_M) and emotion (α_E) separately. For political perspective in news media, $\alpha_M = 0.7$, $\alpha_E = 0.8$, and models were trained for 20 epochs, with early stopping. Within political affiliation prediction, $\alpha_M = \alpha_E = 0.9$ and the first 5 epochs are for auxiliary warm-up. The models were trained for 20 epochs total. For the framing task $\alpha_M = \alpha_E = 0.5$, with 5 epochs of auxiliary warm-up for metaphor. Training lasted 10 epochs at most, with early stopping.

We average results over 10 random seeds. We perform significance testing using an approximated permutation test and 10 thousand permutations.

n	Document Piece	Gold Label	MTL, Metaphor	STL
1	...the anger simmering just below the surface in the U. S. is beginning to boil over.	Right	Right	Centre
2	...and DNA evidence does not match. What once was considered an airtight case, Devine said, has evaporated into nothing	Fairness and Equality	Fairness and Equality	Leg., Constit., Jurisdiction
3	...border security long have been sticking points in the immigration debate. Bowing to those concerns, Presidents Bush ...	Security & Defense	Security & Defense	Political

Table 1: Political perspective (1) and framing (2, 3) examples of metaphor MTL improving over STL. Underlined are words predicted as metaphorical.

	Anger	Anticipation	Disgust	Fear	Joy	Love	Optimism	Pessimism	Sadness	Surprise	Trust
Democrat	34.0%	42.9%	42.2%	23.1%	61.9%	73.6%	54.0%	82.5%	76.4%	75.4%	41.6%
Republican	66.0%	57.1%	57.8%	76.9%	38.1%	26.4%	46.0%	17.5%	23.6%	24.6%	58.4%

Table 2: Proportion of posts predicted for each emotion, using the best-performing MTL emotion model.

	Framing	Affiliation	Perspective
Li and Goldwasser (2019)			
- HLSTM (text-based)	-	-	.746
- GCN-HLSTM (using social information)			.917
Ji and Smith (2017)			
- Tree-RNN[5]	.584	-	-
Khanehzar et al. (2019)			
- RoBERTa [6]	.658	-	-
STL	.707	.794	.848
MTL, Metaphor	**.716**	**.805**	**.854**
MTL, Emotion	.708	**.802**	**.860**

Table 3: Accuracy scores for the main political tasks. Significance compared to STL is bolded ($p < 0.05$).

5.2 Results

Table 3 summarises our results. For political perspective detection in news articles, the STL model improves over the text-based method of Li and Goldwasser (2019). This illustrates that RoBERTa provides an enhanced document encoding for predicting political perspectives. Both MTL setups significantly improved over the STL model. Joint learning with emotion proved most beneficial and outperformed the metaphor detection setup, significantly. For political affiliation prediction, both MTL setups improve over STL significantly, although there is no significant difference between them. In case of the framing task, joint learning with metaphor significantly outperformed STL. MTL using emotion, on the other hand, yielded results on par with STL. Table 4 presents a breakdown of the performance per class in each task.

[5]Ji and Smith (2017) use only a subset (Immigration Policy) of the data and an older version of the dataset, which complicates direct performance comparisons.

[6]Khanehzar et al. (2019) report results on the Immigration Policy subset of the data only.

	STL	MTL, Metaphor	MTL, Emotion
Political Perspective			
- Center	.874	.879	.885
- Left	.860	.863	.871
- Right	.774	.784	.798
Political Affiliation			
- Democrat	.788	.806	.799
- Republican	.802	.805	.800
Framing			
- Economic	.747	.759	.758
- Capacity & Resources	.601	.604	.602
- Morality	.646	.662	.648
- Fairness & Equality	.502	.527	.511
- Crime & Punishment	.719	.721	.717
- Security & Defense	.554	.577	.560
- Health & Safety	.683	.694	.684
- Quality of Life	.572	.554	.556
- Cultural Identity	.690	.703	.695
- Public Sentiment	.670	.678	.675
- Political	.808	.815	.812
- Legality, Constitutionality & Jurisdiction	.787	.795	.784
- Policy Prescription & Evaluation	.525	.538	.530
- External Regulation & Reputation	.695	.675	.681

Table 4: Average F1 for each class and task.

6 Discussion

Political Perspective and Affiliation For the political perspective detection task, the performance improvements of MTL models stem mostly from improved predictions for the *right-wing* class. Example 1 of Table 1 presents an emotive article snippet containing the metaphors of "boil over" and "simmering anger", for which joint learning with metaphor corrected the STL prediction.

For political affiliation, improvements from auxiliary tasks are due to a more accurate identification of the class of *democrats*. According to Pliskin et al.

(2014), liberals are more susceptible to emotions, which could in part explain this result. Appendix A.2.2, Figure 2 visualises the performance across the political spectrum, from which we infer that politicians at the centre are harder to distinguish, and those on the left are better identified by our MTL models. We explored the emotions predicted by the MTL model in politicians' posts, as shown in Table 2. We found that emotions typically associated with conservative rhetoric – e.g. anger, disgust or fear (Jost et al., 2003) – were more frequent in republicans' posts. On the contrary, emotions associated with the liberal rhetoric – e.g. love (Lakoff, 2002) or sadness (Steiger et al., 2019) – are more often predicted for democrats' posts. Appendix A.2.2, Table 7 contains example posts where joint learning using emotion corrected the STL setup.

Framing In case of the framing task, MTL with metaphor prediction yielded the largest improvements for the frames of *security and defence*, *morality* and *fairness and equality*, particularly in articles on the metaphor-rich topics of *immigration*, *gun-control* and *death penalty*. We automatically annotated metaphorical expressions in these articles to conduct a qualitative analysis. We observed that correct identification of linguistic metaphors often accompanies correct frame classification by the MTL model. Examples of such cases are shown in Table 1. In Example 2, metaphors such as "airtight case" and "evaporated", aided the model to identify the *fairness and equality* framing within the topic of *death penalty*. Similarly, presenting border security in Example 3 as a "*sticking point* in the immigration debate" improved the classification of the *security and defense* framing of an article on the topic of *immigration*. Appendix A.2.1, Table 6 presents detailed results per policy issue. The results for the immigration policy subset can be compared to those from Ji and Smith (2017) and Khanehzar et al. (2019).

7 Conclusion

In this paper, we introduced the first joint models of metaphor, emotion and political rhetoric. We considered predicting the political perspective of news, the party affiliation of politicians and the framing dimensions of policy issues. MTL using metaphor detection resulted in significant performance improvements across all three tasks. This finding emphasises the prevalence of metaphor in political discourse and its importance for the identification of framing strategies. Joint learning with emotion yielded significant performance improvements for the political perspective and affiliation tasks, which suggests that the use of emotion is an important political tool, aiming to influence public opinion. Future research may explore further tasks such as emotion and misinformation detection, which social scientists have found to be inter-related, and deploy more advanced MTL techniques, such as soft parameter sharing. Our code is publicly available at github.com/LittlePea13/mtl_political_discourse.

Acknowledgments

We thank Chang Li and Dan Goldwasser for sharing their news dataset on political perspective detection (Li and Goldwasser, 2019); as well as Dallas Card, Yangfeng Ji and Noah A. Smith for sharing their Media Frames Corpus (Card et al., 2015).

This research was funded by the H2020 project *Democratic Efficacy and the Varieties of Populism in Europe* (DEMOS) under H2020-EU.3.6.1.1. and H2020-EU.3.6.1.2. (grant agreement ID: 822590).

References

Muhammad Abdul-Mageed and Lyle Ungar. 2017. EmoNet: Fine-grained emotion detection with gated recurrent neural networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL) (Volume 1: Long Papers)*, pages 718–728.

Cecilia Ovesdotter Alm, Dan Roth, and Richard Sproat. 2005. Emotions from text: Machine learning for text-based emotion prediction. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing (HLT-EMNLP)*, page 579–586.

Beata Beigman Klebanov, Chee Wee Leong, E. Dario Gutierrez, Ekaterina Shutova, and Michael Flor. 2016. Semantic classifications for detection of verb metaphors. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL) (Volume 2: Short Papers)*, pages 101–106.

Lori D. Bougher. 2012. The case for metaphor in political reasoning and cognition. *Political Psychology*, 33(1):145–163.

Luana Bulat, Stephen Clark, and Ekaterina Shutova. 2017. Modelling metaphor with attribute-based semantics. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics (EACL) (Volume 2: Short Papers)*, pages 523–528.

Dallas Card, Amber E. Boydstun, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. The media

frames corpus: Annotations of frames across issues. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (ACL-IJCNLP) (Volume 2: Short Papers)*, pages 438–444.

Jonathan Charteris-Black. 2009. *Metaphor and Political Communication*, pages 97–115. Palgrave Macmillan UK, London.

Francesca MM Citron and Adele E Goldberg. 2014. Metaphorical sentences are more emotionally engaging than their literal counterparts. *Journal of cognitive neuroscience*, 26(11):2585–2595.

Michael D. Conover, Bruno Goncalves, Jacob Ratkiewicz, Alessandro Flammini, and Filippo Menczer. 2011. Predicting the political alignment of twitter users. In *2011 IEEE Third Int'l Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third Int'l Conference on Social Computing*, pages 192–199.

Taner Danisman and Adil Alpkocak. 2008. Feeler: Emotion classification of text using vector space model. In *AISB 2008 Convention Communication, Interaction and Social Intelligence*, volume 1, page 53.

Verna Dankers, Marek Rei, Martha Lewis, and Ekaterina Shutova. 2019. Modelling the interplay of metaphor and emotion through multitask learning. 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 2218–2229.

Derek Edwards. 1999. Emotion discourse. *Culture & Psychology*, 5(3):271–291.

Paul Ekman. 1992. An Argument for Basic Emotions. *Cognition and Emotion*.

Robert M. Entman. 1993. Framing: Toward clarification of a fractured paradigm. *Journal of Communication*, 43(4):51–58.

Federica Ferrari. 2007. Metaphor at work in the analysis of political discourse: investigating a 'preventive war' persuasion strategy. *Discourse & Society*, 18(5):603–625.

Vladimir Figar. 2014. Emotional appeal of conceptual metaphors of conflict in the political discourse of daily newspapers. *Facta Universitatis: Series Linguistics and LIterature*, 12:43–61.

Ge Gao, Eunsol Choi, Yejin Choi, and Luke Zettlemoyer. 2018. Neural metaphor detection in context. *CoRR*, abs/1808.09653.

E. Dario Gutiérrez, Ekaterina Shutova, Tyler Marghetis, and Benjamin Bergen. 2016. Literal and metaphorical senses in compositional distributional semantic models. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL) (Volume 1: Long Papers)*.

Mareike Hartmann, Tallulah Jansen, Isabelle Augenstein, and Anders Søgaard. 2019. Issue framing in online discussion fora. *CoRR*, abs/1904.03969.

Mohit Iyyer, Peter Enns, Jordan Boyd-Graber, and Philip Resnik. 2014. Political ideology detection using recursive neural networks. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL) (Volume 1: Long Papers)*, pages 1113–1122.

Yangfeng Ji and Noah A. Smith. 2017. Neural discourse structure for text categorization. *CoRR*, abs/1702.01829.

Kristen Johnson, Di Jin, and Dan Goldwasser. 2017. Leveraging behavioral and social information for weakly supervised collective classification of political discourse on twitter. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 741–752, Vancouver, Canada. Association for Computational Linguistics.

John Jost, Jack Glaser, Arie Kruglanski, and Frank Sulloway. 2003. Political conservatism as motivated social cognition. *Psychological bulletin*, 129:339–75.

Shima Khanehzar, A. Turpin, and Gosia Mikolajczak. 2019. Modeling political framing across policy issues and contexts. In *ALTA*.

Johannes Kiesel, Maria Mestre, Rishabh Shukla, Emmanuel Vincent, Payam Adineh, David Corney, Benno Stein, and Martin Potthast. 2019. SemEval-2019 Task 4: Hyperpartisan news detection. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 829–839.

Eliyahu Kiperwasser and Miguel Ballesteros. 2018. Scheduled multi-task learning: From syntax to translation. *CoRR*, abs/1804.08915.

Zoltán Kövecses. 2002. *Metaphor: A Practical Introduction*. Oxford University Press, Oxford.

G. Lakoff. 2002. *Moral Politics: How Liberals and Conservatives Think, Second Edition*. University of Chicago Press.

George Lakoff. 1991. Metaphor and war: The metaphor system used to justify war in the gulf. *Peace Research*, 23(2/3):25–32.

George Lakoff and Elisabeth Wehling. 2012. *The Little Blue Book: The Essential Guide to Thinking and Talking Democratic*. Free Press, New York.

Sophie Lecheler, Linda Bos, and Rens Vliegenthart. 2015. The mediating role of emotions: News framing effects on opinions about immigration. *Journalism & Mass Communication Quarterly*, 92(4):812–838.

Julia Lee Cunningham, Yunkyu Sohn, and James Fowler. 2013. Emotion regulation as the foundation of political attitudes: Does reappraisal decrease support for conservative policies? *PloS one*, 8:e83143.

Geoffrey Leech. 1992. 100 million words of english: The british national corpus (bnc).

Chang Li and Dan Goldwasser. 2019. Encoding Social Information with Graph Convolutional Networks forPolitical Perspective Detection in News Media.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692.

Will Lowe, Kenneth Benoit, Slava Mikhaylov, and Michael Laver. 2011. Scaling policy preferences from coded political texts. *Legislative Studies Quarterly*, 36(1):123–155.

Saif Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. SemEval-2018 task 1: Affect in tweets. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 1–17.

Saif Mohammad, Ekaterina Shutova, and Peter Turney. 2016. Metaphor as a medium for emotion: An empirical study. In *Proceedings of the Fifth Joint Conference on Lexical and Computational Semantics*, pages 23–33.

Burt L. Monroe, Michael P. Colaresi, and Kevin M. Quinn. 2008. Fightin' words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political Analysis*, 16(4):372–403.

Rachel R. Mourão and Craig T. Robertson. 2019. Fake news as discursive integration: An analysis of sites that publish false, misleading, hyperpartisan and sensational information. *Journalism Studies*, 20(14):2077–2095.

Andreas Musolff. 2004. *Metaphor and Political Discourse*. Palgrave Macmillan, London.

Marco Pennacchiotti and Ana-Maria Popescu. 2011. Democrats, republicans and starbucks afficionados: user classification in twitter. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 430–438.

Ruthie Pliskin, Daniel Bar-Tal, Gal Sheppes, and Eran Halperin. 2014. Are leftists more emotion-driven than rightists? the interactive influence of ideology and emotions on support for policies. *Personality and Social Psychology Bulletin*, 40(12):1681–1697. PMID: 25381287.

Daniel Preoţiuc-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 (ACL) (Volume 1: Long Papers)*, pages 729–740.

Marek Rei, Luana Bulat, Douwe Kiela, and Ekaterina Shutova. 2017. Grasping the Finer Point: A Supervised Similarity Network for Metaphor Detection. In *Proceedings of EMNLP 2017*.

E. Shutova, L. Sun, and A. Korhonen. 2010. Metaphor identification using verb and noun clustering. In *Coling*.

G.J. Steen, A.G. Dorst, J.B. Herrmann, A.A. Kaal, T. Krennmayr, and T. Pasma. 2010. *A method for linguistic metaphor identification. From MIP to MIPVU.* Number 14 in Converging Evidence in Language and Communication Research. John Benjamins.

Russell L. Steiger, Christine Reyna, Geoffrey Wetherell, and Gabrielle Iverson. 2019. Contempt of congress: Do liberals and conservatives harbor equivalent negative emotional biases towards ideologically congruent vs. incongruent politicians at the level of individual emotions? *Journal of Social and Political Psychology*, 7:100–123.

Deborah Stone. 1988. *Policy paradox and political reason*. Addison-Wesley Longman.

Tomek Strzalkowski, George Aaron Broadwell, Sarah Taylor, Laurie Feldman, Samira Shaikh, Ting Liu, Boris Yamrom, Kit Cho, Umit Boz, Ignacio Cases, and Kyle Elliot. 2013. Robust extraction of metaphor from novel data. In *Proceedings of the First Workshop on Metaphor in NLP*, pages 67–76.

Matt Thomas, Bo Pang, and Lillian Lee. 2006. Get out the vote: Determining support or opposition from Congressional floor-debate transcripts. In *COLING/ACL 2006 - EMNLP 2006: 2006 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*.

Yulia Tsvetkov, Leonid Boytsov, Anatole Gershman, Eric Nyberg, and Chris Dyer. 2014. Metaphor detection with cross-lingual model transfer. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 248–258.

P. D. Turney, Y. Neuman, D. Assaf, and Y. Cohen. 2011. Literal and metaphorical sense identification through concrete and abstract context. In *EMNLP*.

Rob Voigt, David Jurgens, Vinodkumar Prabhakaran, Dan Jurafsky, and Yulia Tsvetkov. 2018. RtGender: A corpus for studying differential responses to gender. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC)*.

Brian E. Weeks. 2015. Emotions, partisanship, and misperceptions: How anger and anxiety moderate

the effect of partisan bias on susceptibility to political misinformation. *Journal of Communication*, 65(4):699–719.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2019. Huggingface's transformers: State-of-the-art natural language processing.

C. Yang, K. H. Lin, and H. Chen. 2007. Emotion classification using web blog corpora. In *IEEE/WIC/ACM International Conference on Web Intelligence (WI'07)*, pages 275–278.

Da Yin, Tao Meng, and Kai-Wei Chang. 2020. SentiBERT: A transferable transformer-based architecture for compositional sentiment semantics. *CoRR*, abs/2005.04114.

A Supplemental Material

A.1 Experimental Setup

Our implementation is in Pytorch and uses Huggingface's Transformers library[7] to load the pretrained models and perform finetuning. Some code for training models was adapted from the utils_nlp[8] library. The data splits and code were attached with our submission.

In order to train the model, a cluster with 4 x Titan RTX, 24 GB GDDR6 GPU with an Intel® Xeon® 2.30 GHz CPU was used. Each model was trained in under two hours. STL models were trained in half of the time it took to train MTL models. We finetuned the pretrained RoBERTa model and do not account for the time it took to pretrain RoBERTa. RoBERTa itself has 125 million parameters and our task-specific layers added around 5 million parameters, with some variance per task, making a total of 130 million parameters.

We experimented with multiple α values at intervals of 0.1. To estimate the warm-up period for scheduled learning, we experimented with 3, 4 or 5 epochs. For the political affiliation task, dropout probabilities of 0.1, 0.2 and 0.3 were experimented with. The final hyperparameter setups were chosen through manual tuning based on the accuracy scores on the validation sets. Hyperparameters that were shared between MTL and STL for the same main task were selected based on the performance in the STL setup. The validation accuracy socres are listed in Table 5.

	Perspective	Affiliation	Framing
STL	.832	.804	.699
MTL, Metaphor	.835	.804	.703
MTL, Emotion	.838	.811	.704

Table 5: Accuracy validation scores for the main political tasks.

A.2 Results

A.2.1 Framing

	STL	MTL, Metaphor	MTL, Emotion
Immigration	0.689	0.700	0.686
Tobacco	0.718	0.721	0.717
Death Penalty	0.690	0.704	0.690
Gun Control	0.704	0.717	0.711
Same Sex Marriage	0.744	0.741	0.745

Table 6: Average accuracy values across different policies for Framing.

A.2.2 Political Affiliation

[7] https://github.com/huggingface/transformers

[8] https://github.com/microsoft/nlp-recipes/tree/master/utils_nlp

Emotions	Facebook Post	Gold Label	emotions MTL	STL
Anger, Disgust and Fear	Last week, I held a Congress on Your Corner event in Frankfort. Monica was upset by the recent deal between the United States, our global partners, and Iran. The deal provides $7 billion in sanction relief in exchange for Iran limiting, but not halting, its nuclear activities. I am skeptical of this deal. In the words of my friend Eric Cantor, I believe we must distrust and verify in this case. I believe it is imperative that we stand with Israel against the very dangerous threat posed by Iran's nuclear activities. I do not believe that Iran has given us any reason to trust that it will not continue pursuing nuclear weapons.	Republican	Republican	Democrat
Love, Joy and Optimism	I'll be spending most of my day tomorrow opposing Paul Ryan's cuts-only budget in committee. In the name of deficit reduction, Mr. Ryan is once again proposing to eliminate one of the few pieces of good news we have in reducing healthcare costs that are driving the deficits: Obamacare (aka, the Affordable Care Act). We should be expanding its reforms, not trying to repeal them. For example, the CBO estimates that adding a public plan option to the health insurance exchanges would save another $88 billion and that the plan would have premiums 5-7% lower than private plans, which would increase competition in the marketplace and result in substantial savings for individuals, families, and employers purchasing health insurance through an exchange.	Democrat	Democrat	Republican

Table 7: Examples from the test set for the Political Affiliation task where emotion detection MTL improved the predictions over STL.

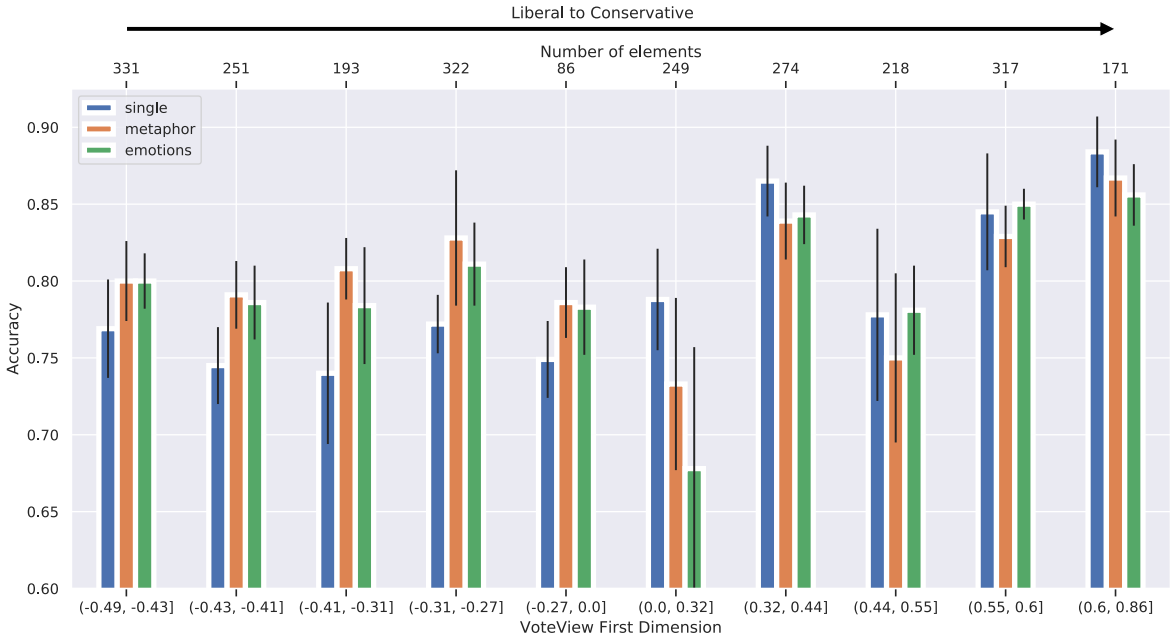

Figure 2: Average performance accross the political spectrum for the Political Affiliation task. Dimension taken from Voteview.

.

A.3 Lists

A.3.1 List of Frames in Card et al. (2015) dataset

- **Economic**: costs, benefits, or other financial implications
- **Capacity** and resources: availability of physical, human or financial resources, and capacity of current systems
- **Morality**: religious or ethical implications
- **Fairness and equality**: balance or distribution of rights, responsibilities, and resources
- **Legality, constitutionality and jurisprudence**: rights, freedoms, and authority of individuals, corporations, and government
- **Policy prescription and evaluation**: discussion of specific policies aimed at addressing problems
- **Crime and punishment**: effectiveness and implications of laws and their enforcement
- **Security and defense**: threats to welfare of the individual, community, or nation
- **Health and safety**: health care, sanitation, public safety
- **Quality of life**: threats and opportunities for the individual's wealth, happiness, and well-being
- **Cultural identity**: traditions, customs, or values of a social group in relation to a policy issue
- **Public opinion**: attitudes and opinions of the general public, including polling and demographics
- **Political**: considerations related to politics and politicians, including lobbying, elections, and attempts to sway voters
- **External regulation and reputation**: international reputation or foreign policy of the U.S.
- **Other**: any coherent group of frames not covered by the above categories

A.3.2 List of emotions in Mohammad et al. (2018) dataset

- **Anger** (also includes annoyance, rage)
- **Anticipation** (also includes interest, vigilance)
- **Disgust** (also includes disinterest, dislike, loathing)
- **Fear** (also includes apprehension, anxiety, terror)
- **Joy** (also includes serenity, ecstasy)
- **Love** (also includes affection)
- **Optimism** (also includes hopefulness, confidence)
- **Pessimism** (also includes cynicism, no confidence)
- **Sadness** (also includes pensiveness, grief)
- **Surprise** (also includes distraction, amazement)
- **Trust** (also includes acceptance, liking, admiration)