

Power of Confidence: How Poll Scores Impact Topic Dynamics in Political Debates

Vinodkumar Prabhakaran

Dept. of Computer Science
Columbia University
New York, NY

vinod@cs.columbia.edu

Ashima Arora

Dept. of Computer Science
Columbia University
New York, NY

aa3470@columbia.edu

Owen Rambow

CCLS
Columbia University
New York, NY

rambow@ccls.columbia.edu

Abstract

In this paper, we investigate how topic dynamics during the course of an interaction correlate with the power differences between its participants. We perform this study on the US presidential debates and show that a candidate’s power, modeled after their poll scores, affects how often he/she attempts to shift topics and whether he/she succeeds. We ensure the validity of topic shifts by confirming, through a simple but effective method, that the turns that shift topics provide substantive topical content to the interaction.

1 Introduction

Analyzing political speech has gathered great interest within the NLP community. Researchers have analyzed political text to identify markers of persuasion (Guerini et al., 2008), predict voting patterns (Thomas et al., 2006; Gerrish and Blei, 2011), and detect ideological positions (Sim et al., 2013). Studies have also looked into how personal attributes of political personalities such as charisma, confidence and power affect how they interact (Rosenberg and Hirschberg, 2009; Prabhakaran et al., 2013b). Our work belongs to this genre of studies. We analyze how a presidential candidate’s power, modeled after his/her relative poll standings, affect the dynamics of topic shifts during the course of a presidential debate.

2 Motivation

In early work on correlating personal attributes to political speech, Rosenberg and Hirschberg (2009) analyzed speech transcripts in the context of 2004 Democratic presidential primary elections, to identify prosodic and lexico-syntactic cues that signal charisma of political personalities.

More recently, Prabhakaran et al. (2013a) introduced the notion of power an election candidate has at a certain point in the election campaign, modeled after the confidence that stems from their recent poll standings. They analyzed the 2012 Republican presidential primary debates and found that the candidate’s power at the time of a debate impacts the structure of interactions (e.g., frequency of turns and interruption patterns). They followed up their study with an automatic ranker to identify leading candidates based on the interaction within a debate (Prabhakaran et al., 2013b).

One of the interesting findings by Prabhakaran et al. (2013a) was that candidates’ power correlates with the distribution of topics they speak about in the debates. They found that when candidates have more power, they speak significantly more about certain topics (e.g., *economy*) and less about certain other topics (e.g., *energy*). However, these findings relate to the specific election cycle they analyzed and will not carry over to all political debates in general. A topical dimension with broader relevance is how topics change during the course of an interaction (e.g., who introduces more topics, who attempts to shift topics etc.). For instance, Nguyen et al. (2013) found that topic shifts within an interaction are correlated with the role a participant plays in it (e.g., being a moderator). They also analyzed US presidential debates, but with the objective of validating a topic segmentation method they proposed earlier (Nguyen et al., 2012). They do not study the topic shifting tendencies among the candidates in relation to their power differences.

In this paper, we bring these two ideas together. We analyze the 2012 Republican presidential debates, modeling the power of a candidate based on poll scores as proposed by Prabhakaran et al. (2013a) and investigate various features that capture the topical dynamics in the debates. We show that the power affects how often candidates at-

Turn #	Speaker	Turn Text	Substantive?
223	PAWLENTY (C)	I support a constitutional amendment to define marriage between a man and woman. I was the co-author of the state – a law in Minnesota to define it and now we have courts jumping over this.	[S]
224	KING (M)	OK. Let’s just go through this.	[NS]
225	PAUL (C)	The federal government shouldn’t be involved. I wouldn’t support an amendment. [...] I don’t think government should give us a license to get married. It should be in the church.	[S]
226	KING (M)	Governor Romney, constitutional amendment or state decision?	[NS]
227	ROMNEY (C)	Constitutional.	[NS]
228	KING (M)	Mr. Speaker?	[NS]
229	GINGRICH (C)	Well, I helped author the Defense of Marriage Act which the Obama administration should be frankly protecting in court. [...]	[S]
		[...]	
235	CAIN (C)	If I had my druthers, I never would have overturned “don’t ask/don’t tell” in the first place. [...] Our men and women have too many other things to be concerned about rather than have to deal with that as a distraction.	[S]
		[...]	
240	KING (M)	Leave it in place, [...] or overturn it?	[S]
241	ROMNEY (C)	Well, one, we ought to be talking about the economy and jobs. But given the fact you’re insistent, the – the answer is, I believe that “don’t ask/don’t tell” should have been kept in place until conflict was over.	[S]

Table 1: Excerpt from Goffstown, NH debate (06/13/11), discussing marriage equality and the “Don’t Ask/Don’t Tell” policy [S]/ [NS] denote substantiveness of turns

tempt to shift topics and whether they succeed in it or not. In order to correctly model topic shifts, we ensure that the shifts happen in turns that contribute substantial topical content to the interaction. We introduce the notion of a “non-substantial turn”, and use a simple, but effective method to automatically identify non-substantial turns. This allows us to identify different topic segments within the interaction, while permitting (and capturing) interruptions within those segments. We will compare the segments that we obtain with those by Nguyen et al. (2012) in future work.

3 Domain and Data

We use the same corpus as Prabhakaran et al. (2013b). The corpus contains manual transcripts of 20 debates held between May 2011 and February 2012 as part of the 2012 Republican presidential primaries. The transcripts are obtained from The American Presidency Project.¹ Each turn is clearly demarcated in the transcripts and their speakers are identified. The turns in the corpus are preprocessed using the Stanford CoreNLP package to perform basic NLP steps such as tokenization, sentence segmentation, parts-of-speech tagging and lemmatization. We show an excerpt

¹<http://www.presidency.ucsb.edu/debates.php>

from one of the debates in Table 1. This segment of the debate discusses marriage equality followed by the overturning of the “Don’t Ask/Don’t Tell” policy prohibiting openly gay, lesbian, or bisexual persons from US military service.

Prabhakaran et al. (2013b) added each candidate’s power at the time of each debate to the corpus, computed based on their relative standing in recent public polls. We refer the reader to (Prabhakaran et al., 2013b) for the detailed description of how the relative standings in national and state-level polls from various sources are aggregated to obtain candidates’ power. The poll numbers capture how successful candidates are in convincing the electorate of their candidature, which in turn affects their confidence within the debates. These debates serve as a rich domain to explore manifestations of power since they are a medium through which candidates pursue and maintain power over other candidates.

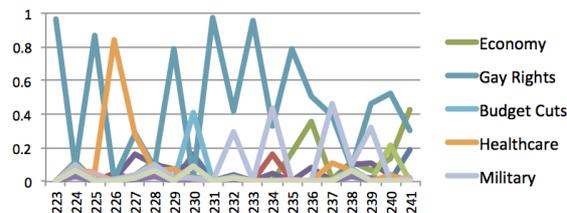
4 Modeling Topics

Prabhakaran et al. (2013a) model topics in the debates using Latent Dirichlet Allocation (LDA), assigning topic probabilities to each turn. The number of topics was set to be 15 and the topic that was assigned the highest probability for a turn was cho-

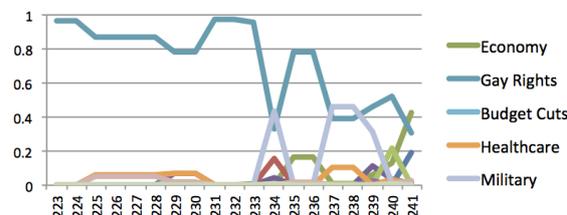
sen as its topic. Assigning topics to each turn in this manner, however, is problematic. Not all turns by themselves contribute to the conversational topics in an interaction. A large number of turns, especially by the moderator, manage the conversation rather than contribute content to it. These include turns redirecting questions to specific candidates (e.g., turns 224, 226 and 228 in Table 1) as well as moderator interruptions (e.g., “Quickly.”, “We have to save time”). Furthermore, some other turns address a topic only when considered together with preceding turns, but not when read in isolation. These include turns that are short one-word answers (e.g., turn 227) and turns that are uninterpretable without resolving anaphora (e.g., “That’s right”). While these turns are substantive to human readers, topic modeling approaches such as LDA cannot assign them topics correctly because of their terseness.

We define the turns that do not, in isolation, contribute substantially to the conversational topics as **non-substantive** turns. In order to obtain a gold standard for non-substantivity, two of the authors manually annotated each turn in one entire debate (dated 06/13/11) as either *substantive* (*S*) or *non-substantive* (*NS*). The annotators were instructed not to consider the identity of the speaker or the context of the turn (preceding/following turns) in making their assessment. We obtained a high inter-annotator agreement (observed agreement = 89.3%; Kappa = .76). We took the assessments by one of the annotators as the gold standard, in which 108 (31.5%) of the 343 turns were identified as *non-substantive*. We show the *S* vs. *NS* assessments for each turn in column 4 of Table 1.

Figure 1a shows the line graph of topic probabilities assigned by LDA to the sequence of turns in Table 1. As the graph shows, non-substantive turns are assigned spurious topic probabilities by LDA. For example, turn 224 by KING (“OK. Lets just go through this.”) was assigned small probabilities for all topics; the highest of which was *economy* (probability of 0.12). This error is problematic when modeling topic shifts, since this turn and the next one by PAUL would have been incorrectly identified as shifts in topic from their corresponding previous turns. Instead, if we assume that the non-substantive turns follow the same topic probabilities as the most recent substantive turn, we obtain the line graph shown in Figure 1b. This topic assignment captures the topic dynam-



(a) Topic Probabilities assigned by LDA



(b) Topic Probabilities after ignoring non-substantive turns

Figure 1: Line graphs of topic probabilities for turns in Table 1 (legend shows only the top 5 topics in this segment)

ics in the segment more accurately. It identifies *Gay Rights* as the predominant topic until turn 234 followed by a mix of *Gay Rights* and *Military* as topics while discussing the “Don’t Ask/Don’t Tell” policy. It also captures the attempt by ROMNEY in turn 242 to shift the topic to *Economy*.

4.1 Identifying Non-substantive Turns

In order to automatically detect non-substantive turns, we investigate a few alternatives. A simple observation is that many of the *NS* turns such as redirections of questions or short responses have only a few words. We tried a word count threshold based method (**WC_Thresh**) where we assign a turn to be *NS* if the number of tokens (words) in the turn is less than a threshold. Another intuition is that for a non-substantive turn, it would be hard for the LDA to assign topics and hence all topics will get almost equal probabilities assigned. In order to capture this, we used a method based on a standard deviation threshold (**SD_Thresh**), where we assign a turn to be *NS* if the standard deviation of that turn’s topic probabilities is below a threshold. We also used a combination system where we tag a turn to be *NS* if either system tags it to be. We tuned for the value of the thresholds and the best performances obtained for each case are shown in Table 2. We obtained the best results for the WC_Thresh method with a threshold of 28 words, while for SD_Thresh the optimal threshold is .13 (almost twice the mean).

Method	Accuracy (%)	F-measure
WC_Thresh	82.6	73.7
SD_Thresh	76.2	64.7
WC_Thresh + SD_Thresh	76.8	70.4

Table 2: Accuracy and F-measure of different methods to identify non-substantive turns

4.2 Topic Assignments

We first ran the LDA at a turn-level for all debates, keeping the number of topics to be 15, and selected the best model after 2000 iterations. Then, we ran the WC_Thresh method described above to detect *NS* turns. For all *NS* turns, we replace the topic probabilities assigned by LDA with the last substantive turn’s topic probabilities. Note that an *S* turn coming after one or more *NS* turns could still be of the same topic as the last *S* turn, i.e., non-substantivity of a turn is agnostic to whether the topic changes after that or not. A topic shift (or attempt) happens only when LDA assigns a different topic to a substantive turn.

5 Topical Dimensions

We now describe various features we use to capture the topical dynamics within each debate, with respect to each candidate. When we compute a feature value, we use the topic probabilities assigned to each turn as described in the previous section. For some features we only use the topic with the highest probability, while for some others, we use the probabilities assigned to all topics. We consider features along four dimensions which we describe in detail below.

5.1 Topic Shift Patterns

We build various features to capture how often a candidate stays on the topic being discussed. We say a candidate attempted to shift the topic in a turn if the topic assigned to that turn differs from the topic of the previous (substantive) turn. We use a feature to count the number of times a candidate attempts to shift topics within a debate (**TS_Attempt#**) and a version of that feature normalized over the total number of turns (**TS_Attempt#^N**). We also use a variation of these features which considers only the instances of topic shift attempts by the candidates when responding to a question from the moderator (**TS_AttemptAfterMod#** and

TS_AttemptAfterMod#^N). We also compute a softer notion of topic shift where we measure the average Euclidean distance between topic probabilities of each of the candidate turns and turns prior to them (**EuclideanDist**). This feature in essence captures whether the candidate stayed on topic, even if he/she did not completely switch topics in a turn.

5.2 Topic Shift Sustenance Patterns

We use a feature to capture the average number of turns for which topic shifts by a candidate was sustained (**TS_SustTurns**). However, as discussed in Section 4, the turns vary greatly in terms of length. A more sensible measure is the time period for which a topic shift was sustained. We approximate the time by the number of word tokens and compute the average number of tokens in the turns that topic shifts by a candidate were sustained (**TS_SustTime**).

5.3 Topic Shift Success Patterns

We define a topic shift to be successful if it was sustained for at least three turns. We compute three features — total number of successful topic shifts by a candidate (**TS_Success#**), that number normalized over the total number of turns by the candidate (**TS_Success#^N**), and the success rate of candidate’s topic shifts (**TS_SuccessRate**)

5.4 Topic Introduction Patterns

We also looked at cases where a candidate introduces a new topic, i.e., shifts to a topic which is entirely new for the debate. We use the number of topics introduced by a candidate as a feature (**TS_Intro#**). We also use features to capture how important those topics were, measured in terms of the number of turns about those topics in the entire debate (**TS_IntroImpTurns**) and the time spent on those topics in the entire debate (**TS_IntroImpTime**).

6 Analysis and Results

We performed a correlation analysis on the features described in the previous section with respect to each candidate against the power he/she had at the time of the debate (based on recent poll scores). Figure 2 shows the Pearson’s product correlation between each topical feature and candidate’s power. Dark bars denote statistically significant ($p < 0.05$) features.

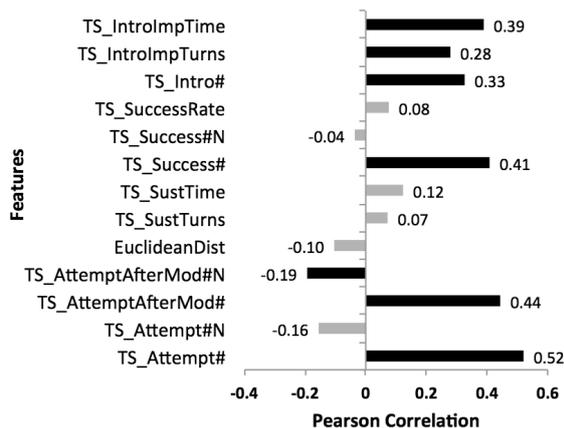


Figure 2: Pearson Correlations for Topical Features

We obtained significant strong positive correlation for `TS_Attempt#` and `TS_AttemptAfterMod#`. However, the normalized measure `TS_Attempt#N` did not have any significant correlation, suggesting that the correlation obtained for `TS_Attempt#` is mostly due to the fact that candidates with more power have more turns, a finding that is already established by Prabhakaran et al. (2013b). However, interestingly, we obtained a weak, but statistically significant, negative correlation for `TS_AttemptAfterMod#N` which suggests that more powerful candidates tend to stay on topic when responding to moderators. We did not obtain any correlation for `EuclideanDist`.

We did not obtain any significant correlations between candidate’s power and their topic shift sustenance features. We obtained significant correlation for topic shift success (`TS_Success#`), modeled based on the sustenance of topic shifts, suggesting that powerful candidates have a higher number of successful topic shifts. However, `TS_SuccessRate` or `TS_Success#N` did not obtain any significant correlation. We also found that powerful candidates are more likely to introduce new topics (`TS_Intro#`) and that the topics they introduce tend to be important (`TS_IntroImpTurns` and `TS_IntroImpTime`).

7 Related Work

Studies in sociolinguistics (e.g., (Ng et al., 1993; Ng et al., 1995)) have explored how dialog structure in interactions relates to power and influence. Reid and Ng (2000) identified that factors such as frequency of contribution, proportion of turns, and number of successful interruptions are important indicators of influence. Within the dialog commu-

nity, researchers have studied notions of control and initiative in dialogs (Walker and Whittaker, 1990; Jordan and Di Eugenio, 1997). Walker and Whittaker (1990) define “control of communication” in terms of whether the discourse participants are providing new, unsolicited information in their utterances. Their notion of control differs from our notion of power; however, the way we model topic shifts is closely related to their utterance level control assignment. Within the NLP community, researchers have studied power and influence in various genres of interactions, such as organizational email threads (Bramsen et al., 2011; Gilbert, 2012; Prabhakaran and Rambow, 2013), online discussion forums (Danescu-Niculescu-Mizil et al., 2012; Biran et al., 2012) and online chat dialogs (Strzalkowski et al., 2012). The correlates analyzed in these studies range from word/phrase patterns, to derivatives of such patterns such as linguistic coordination, to deeper dialogic features such as argumentation and dialog acts. Our work differs from these studies in that we study the correlates of power in topic dynamics. Furthermore, we analyze spoken interactions.

8 Conclusion

We studied the topical dynamics in the 2012 US presidential debates and investigated their correlation with the power differences between candidates. We showed that a candidate’s power, modeled after their poll scores, has significant correlation with how often he/she introduces new topics, attempts to shift topics, and whether they succeed in doing so. In order to ensure the validity of our topic shifts we devised a simple yet effective way to eliminate turns which do not provide substantial topical content to the interaction. Furthermore, this allowed us to identify different topic segments within the interaction. In future work, we will explore how our way of identifying segments compares to other approaches on topic segmentation in interactions (e.g., (Nguyen et al., 2012)).

Acknowledgments

This paper is based upon work supported by the DARPA DEFT Program. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government. We also thank Debanjan Ghosh and several anonymous reviewers for their constructive feedback.

References

- Or Biran, Sara Rosenthal, Jacob Andreas, Kathleen McKeown, and Owen Rambow. 2012. Detecting influencers in written online conversations. In *Proceedings of the Second Workshop on Language in Social Media*, pages 37–45, Montréal, Canada, June. Association for Computational Linguistics.
- Philip Bramsen, Martha Escobar-Molano, Ami Patel, and Rafael Alonso. 2011. Extracting social power relationships from natural language. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 773–782, Portland, Oregon, USA, June. Association for Computational Linguistics.
- Cristian Danescu-Niculescu-Mizil, Lillian Lee, Bo Pang, and Jon Kleinberg. 2012. Echoes of power: language effects and power differences in social interaction. In *Proceedings of the 21st international conference on World Wide Web, WWW '12*, New York, NY, USA. ACM.
- Sean Gerrish and David Blei. 2011. Predicting legislative roll calls from text. In Lise Getoor and Tobias Scheffer, editors, *Proceedings of the 28th International Conference on Machine Learning, ICML '11*, pages 489–496, New York, NY, USA, June. ACM.
- Eric Gilbert. 2012. Phrases that signal workplace hierarchy. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work, CSCW '12*, pages 1037–1046, New York, NY, USA. ACM.
- Marco Guerini, Carlo Strapparava, and Oliviero Stock. 2008. Corps: A corpus of tagged political speeches for persuasive communication processing. *Journal of Information Technology & Politics*, 5(1):19–32.
- Pamela W. Jordan and Barbara Di Eugenio. 1997. Control and initiative in collaborative problem solving dialogues. In *Working Notes of the AAAI Spring Symposium on Computational Models for Mixed Initiative*, pages 81–84.
- Sik Hung Ng, Dean Bell, and Mark Brooke. 1993. Gaining turns and achieving high in influence ranking in small conversational groups. *British Journal of Social Psychology*, page 32, 265–275.
- Sik Hung Ng, Mark Brooke, and Michael Dunne. 1995. Interruption and in influence in discussion groups. *Journal of Language and Social Psychology*, pages 14(4),369–381.
- Viet-An Nguyen, Jordan Boyd-Graber, and Philip Resnik. 2012. Sits: A hierarchical nonparametric model using speaker identity for topic segmentation in multiparty conversations. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 78–87, Jeju Island, Korea, July. Association for Computational Linguistics.
- Viet-An Nguyen, Jordan Boyd-Graber, Philip Resnik, Deborah A. Cai, Jennifer E. Midberry, and Yuanxin Wang. 2013. Modeling topic control to detect influence in conversations using nonparametric topic models. *Machine Learning*, pages 1–41.
- Vinodkumar Prabhakaran and Owen Rambow. 2013. Written dialog and social power: Manifestations of different types of power in dialog behavior. In *Proceedings of the IJCNLP*, pages 216–224, Nagoya, Japan, October. Asian Federation of Natural Language Processing.
- Vinodkumar Prabhakaran, Ajita John, and Dorée D. Seligmann. 2013a. Power dynamics in spoken interactions: a case study on 2012 republican primary debates. In *Proceedings of the 22nd international conference on World Wide Web companion*, pages 99–100. International World Wide Web Conferences Steering Committee.
- Vinodkumar Prabhakaran, Ajita John, and Dorée D. Seligmann. 2013b. Who had the upper hand? ranking participants of interactions based on their relative power. In *Proceedings of the IJCNLP*, pages 365–373, Nagoya, Japan, October. Asian Federation of Natural Language Processing.
- Scott A. Reid and Sik Hung Ng. 2000. Conversation as a resource for in influence: evidence for prototypical arguments and social identification processes. *European Journal of Social Psych.*, pages 30, 83–100.
- Andrew Rosenberg and Julia Hirschberg. 2009. Charisma perception from text and speech. *Speech Communication*, 51(7):640–655.
- Yanchuan Sim, Brice D. L. Acree, Justin H. Gross, and Noah A. Smith. 2013. Measuring ideological proportions in political speeches. In *Proceedings of the 2013 Conference on EMNLP*, pages 91–101, Seattle, Washington, USA, October. Association for Computational Linguistics.
- Tomek Strzalkowski, Samira Shaikh, Ting Liu, George Aaron Broadwell, Jenny Stromer-Galley, Sarah Taylor, Umit Boz, Veena Ravishankar, and Xiaoi Ren. 2012. Modeling leadership and influence in multi-party online discourse. In *Proceedings of COLING*, pages 2535–2552, Mumbai, India, December. The COLING 2012 Organizing Committee.
- Matt Thomas, Bo Pang, and Lillian Lee. 2006. Get out the vote: Determining support or opposition from congressional floor-debate transcripts. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pages 327–335, Sydney, Australia, July. Association for Computational Linguistics.
- Marilyn Walker and Steve Whittaker. 1990. Mixed initiative in dialogue: An investigation into discourse segmentation. In *Proceedings of the 28th annual meeting on Association for Computational Linguistics*, pages 70–78. Association for Computational Linguistics.