

Topical Positioning: A New Method for Predicting Opinion Changes in Conversation

Ching-Sheng Lin¹, Samira Shaikh¹, Jennifer Stromer-Galley^{1,2},
Jennifer Crowley¹, Tomek Strzalkowski^{1,3}, Veena Ravishankar¹

¹State University of New York - University at Albany, NY 12222 USA

²Syracuse University

³Polish Academy of Sciences

clin3@albany.edu, sshaikh@albany.edu, tomek@albany.edu

Abstract

In this paper, we describe a novel approach to automatically detecting and tracking discussion dynamics in Internet social media by focusing on attitude modeling of topics. We characterize each participant's attitude towards topics as Topical Positioning, employ Topical Positioning Map to represent the positions of participants with respect to each other and track attitude shifts over time. We also discuss how we used participants' attitudes towards system-detected meso-topics to reflect their attitudes towards the overall topic of conversation. Our approach can work across different types of social media, such as Twitter discussion and online chat room. In this article, we show results on Twitter data.

1 Introduction

The popularity of social networks and the new kinds of communication they support provides never before available opportunities to examine people behaviors, ideas, and sentiments in various forms of interaction. One of the active research subjects is to automatically identify sentiment, which has been adopted in many different applications such as text summarization and product review. In general, people express their stances and rationalize their thoughts on the topics in social media discussion platform. Moreover, some of them explicitly or implicitly establish strategies to persuade others to embrace his/her belief. For example, in the discussion of the topic "Should the legal drinking age be lowered to 18", the participants who are against it may state their views explicitly and list negative consequences of lowering

drinking age to 18 in an attempt to change opinions of those who appear to support the change. This phenomenon actually involves two research problems which have been of great interest in Natural Language Processing: opinion identification and sociolinguistic modeling of discourse. The first problem can be addressed by traditional opinion analysis that recognizes which position or stance a person is taking for the given topics (Somandaran and Wiebe, 2009). The second part requires modeling the sociolinguistic aspects of interactions between participants to detect more subtle opinion shifts that may be revealed by changes in interpersonal conversational dynamics. In this paper, we bring these two research avenues together and describe a prototype automated system that: (1) discovers each participant's position polarities with respect to various topics in conversation, (2) models how participants' positions change over the course of conversation, and (3) measures the distances between participants' relative positions on all topics. We analyzed discussions on Twitter to construct a set of meso-topics based on the persistence of certain noun phrases and co-referential expressions used by the participants. A meso-topic is any local topic in conversation referred to by a noun phrase and subsequently mentioned again at least 5 times via repetition, pronoun or synonym. Meso-topics do not necessarily represent actual topics of conversations, but certainly are important interactive handles used by the speakers. It is our hypothesis that meso-topics can be effectively used to track and predict polarity changes in speakers' positions towards the overall topic of conversation. Once the meso-topics and their polarities for each participant are determined, we can generate a *topical positioning map (or network)* (TPN) showing relative distances between

participants based on all meso-topics in discourse. Comparing different snapshots of the TPN over time, we can observe how the group’s dynamic changes, i.e., how some participants move closer to one another while others drift apart in the discussion. In particular, we suggest that TPN changes can track and predict participants’ changes of opinion about the overall topic of conversation.

The remainder of this paper is organized as follows. In Section 2, we review related work. In Section 3, we describe the components of the proposed technique and the way they are used to implement the system. In Section 4, we discuss initial empirical studies, including data collection and evaluation. In final section, we present conclusions and some future work.

2 Related Work

While systematic research on opinion tracking and influence in dialogues is a relatively new area of computational linguistics, related research includes automatic opinion mining and sentiments extraction from text (Wiebe et al., 2005; Strapparava and Mihalcea, 2008), speech (Vogt et al., 2008) and social networking sites (Martineau and Finin, 2009). Much of the recent work was focused on automatic analysis of product reviews (books, movies, etc.) and extracting customers’ opinions from them (Hu and Liu, 2004; David and Pinch, 2006; Zhuang et al., 2006). A typical approach is to count the number of ‘opinion’ words within a text window around the product names, possibly augmented with syntactic parsing to get dependencies right. An opinion mining application can extract either full opinion sentences (Philip et al., 2003) or may generate a more structured representation (Hu and Liu, 2004). Another recent application of sentiment analysis is ECO system (Effective Communication Online) (Small et al., 2010) that constructs a model of a community-wide sentiment towards certain common issues discussed in social media, particularly forums and open blogs. This model is then used to assess whether a new post would fit into the targeted community by comparing the sentiment polarities about the concepts in the message and in the model. Potential posters are then guided in ways to shape their communication so that it minimizes the number of conflicting concept sentiments, while still preserving the intended message.

Another related research domain is about modeling the social phenomena in discourse. (Strzalkowski et al., 2010, Broadwell et al., 2012) proposed a two-tier approach that relies on extracting observable linguistic features of conversational text to detect mid-level social behaviors such as Topic Control, Disagreement and Involvement. These social behaviors are then used to infer higher-level social roles such as Leader and Influencer, which may have impact on how other participants’ opinions form and change.

3 System Modules

In this section, we describe a series of modules in our system, which include meso-topic extraction, topical positioning and topical positioning map, and explain how we capture opinion shifts.

3.1 Meso-Topic Extraction

Participants mention many ideas and subjects in dialogue. We call these Local Topics, which are any noun phrases introduced that are subsequently mentioned via repetition, synonym, or pronoun (Strzalkowski et al., 2010) by the same participant or different participants. Some local topics persist for only a couple of turns, others for much longer; some are closely relevant to the overall discussion, while others may appear to be digressions. We identify local topics, their first mentions and subsequent mentions, and track participants who make these mentions. Once local topics have been introduced into the dialogue we track their persistence as topic chains, through repetitions of the noun phrase as well as references via pronouns and the use of synonyms. Topic chains do not have to be continuous, they may contain gaps. The lengths of these gaps are also important to measures for some behaviors. Meso-topics are the most persistent local topics, topics that are widely cited through long stretches of discourse. A selection of meso-topics is closely associated with the task in which the discourse participants are engaged. Short “gaps” in the chain are permitted (up to 10 turns, to accommodate digressions, obscure references, noise, etc.). Meso-topics can be distinguished from the local topics because the participants often make polarized statements about them. We use the Stanford part-of-speech tagger (Klein and Manning, 2003) to automatically detect nouns and noun phrases in dialogue and select those with subsequent men-

tions as local topics using a fairly simple pronoun resolution method based primarily on presence of specific lexical features as well as temporal distance between utterances. Princeton Wordnet (Fellbaum et al., 2006) is consulted to identify synonyms and other related words commonly used in co-references. The local topics that form sufficiently long co-reference chains are designated as meso-topics.

3.2 Topical Positioning

Topical Positioning is defined as the attitude a speaker has towards the meso-topics of discussion. Speakers in a dialogue, when discussing issues, especially ones with some controversy, will establish their attitude on each topic, classified as for, against, or neutral/undecided. In so doing, they establish their positions on the issue or topic, which shapes the agenda of the discussion and also shapes the outcomes or conclusions of the discussion. Characterizing topical positioning allows us to see the speakers who are for, who are against, and who are neutral/undecided on a given topic or issue.

To establish topical positioning, we first identify meso-topics that are present in a discourse. For each utterance made by a speaker on a meso-topic we then establish its polarity, i.e., if this utterance is ‘for’ (positive) or ‘against’ (negative), or neutral on the topic. We distinguish three forms of meso-topic valuation that may be present: (a) *express advocacy/disadvocacy*, when the valuation is applied directly to the topic (e.g., “I’m for Carla”); (b) *supporting/dissenting information*, when the valuation is made indirectly by offering additional information about the topic (e.g., “He’s got experience with youngsters.”); and (c) *express agreement/disagreement* with a polarized statement made by another speaker.

The following measures of Topical Positioning are defined: Topic Polarity Index, which establishes the polarity of a speaker’s attitude towards the topic, and Polarity Strength Index, which measures the magnitude of this attitude.

[*Topic Polarity Index (TPX)*] In order to detect the polarity of Topical Positioning on meso-topic T, we count for each speaker:

- All utterances on T using statements with polarity P applied directly to T using appropriate

adverb or adjective phrases, or when T is a direct object of a verb. Polarities of adjectives and adverbs are taken from the expanded ANEW lexicon (Bradley and Lang, 1999).

- All utterances that offer information with polarity P about topic T.
- All responses to other speakers’ statements with polarity P applied to T. In the Twitter environment (and the like), for now we include a re-tweet in this category.

Given these counts we can calculate TPX for each speaker as a proportion of positive, negative and neutral polarity utterances made by this speaker about topic T. A speaker whose utterances are overwhelmingly positive (80% or more) has a pro-topic position (TPX = +1); a speaker whose utterances are overwhelmingly negative takes an against-topic position (TPX = -1); a speaker whose utterances are largely neutral or whose utterances vary in polarity, has a neutral/undecided position on the topic (TPX = 0).

[*Polarity Strength Index (PSX)*] In addition to the valence of the Topical Positioning, we also wish to calculate its strength. To do so, we calculate the proportion of utterances on the topic made by each speaker to all utterances made about this topic by all speakers in the discourse. Speakers, who make most utterances on the topic relative to other speakers, take a stronger position on this topic. PSX is measured on a 5-point scale corresponding to the quintiles in normal distribution.

Topical Positioning Measure (TPM)

In order to establish the value of Topical Positioning for a given topic we combine the values of TPX*PSX. Topical Positioning takes values between +5 (strongest pro) to 0 (neutral/undecided) to -5 (strongest against). For example, a speaker who makes 25% of all utterances on the topic “Carla” (group mean is 12%) and whose most statements are positive, has the strongest pro Topical Positioning on Carla: +5 (for fifth quintile on the positive side).

3.3 Topical Positioning Map (TPN)

Given the combined values of TPM for each participant in a group, we can calculate distances between the speakers on each meso-topic as well as on all meso-topics in a conversation. For meso-

topics (t_1, \dots, t_N), the distance is calculated using a cosine between speakers' "vectors" ($TPM_{t_1}(A) \dots TPM_{t_N}(A)$) and ($TPM_{t_1}(B) \dots TPM_{t_N}(B)$). Specifically, we use $(1 - \text{Cosine}(V_1, V_2))$ to represent distance between node V_1 and V_2 in the network, where the range becomes 0 to 2.

With the aid of TPN, we can detect the opinion shifts and model the impact of speakers with specific social roles in the group, which in our case is the influencer. An influencer is a group participant who has credibility in the group and introduces ideas that others pick up on or support. An influencer model is generated from mid-level sociolinguistic behaviors, including Topic Control, Disagreement and Involvement (Shaikh et al., 2012). In order to calculate effect of the influencer on a group, we track *changes* in the TPN distances between speakers, and particularly between the influencer and other speakers. We want to know if the other speakers in the group moved closer to or further away from the influencer, who may be promoting a particular position on the overall subject of discussion. Our hypothesis is that other participants will move closer (as a group, though not necessarily individually) to an influential speaker. We may also note that some speakers move closer while others move away, indicating a polarizing effect of an influential speaker. If there is more than one influencer in the group these effects may be still more complex.

4 Data Collection and Experiment

Our initial focus has been on Twitter discussions which enable users to create messages, i.e., "tweets". There are plenty of tweet messages generated all the time and it is reported that Twitter has surpassed 400 million tweets per day. With the Twitter API, it is easy to collect those tweets for research, as the communications are considered public. However, most of data obtained publicly is of limited value due to its complexity, lack of focus, and inability to control for many independent variables. In order to derive reliable models of conversational behavior that fulfill our interests in opinion change, we needed a controlled environment with participants whose initial opinions were known and with conversation reasonably focused on a topic of interest. To do so, we recruited participants for a two-week Twitter debates on a variety of issues, one of the topics was "Should the mini-

um legal drinking age be lowered to 18?" We captured participants' initial positions through surveys before each debate, and their exit positions through surveys after the debate was completed two weeks later. The surveys were designed to collect both the participants' opinions about the overall topic of conversation as well as about the roles they played in it. These data were then compared to the automatically computed TPN changes.

4.1 Data Collection

To obtain a suitable dataset, we conducted two groups of controlled and secured experiments with Twitter users. The experiment was specially designed to ensure that participants stay on topic of discussion and that there was a minority opinion represented in the group. We assigned the same **overall topic** for both groups: "lowering the drinking age from 21 to 18". Before the discussion, the participants completed an 11-question survey to determine their pre-discussion attitudes toward overall topic. One participant with the minority opinion was then asked to act as an influencer in the discussion, i.e., to try to convince as many people as possible to adopt his or her position. After the discussion, the participants were asked the same 11 questions to determine if their positions have changed. All 11 questions probed various aspects of the overall topic, thus providing a reliable measure of participant's opinion. All responses were on a 7-point scale from "strongly agree" to "strongly disagree". The orientation of individual questions vs. the overall topic was varied to make sure that the participants did not mechanically fill their responses. Some of the questions were:

(1) *Lowering the drinking age to 18 would make alcohol less of a taboo, making alcohol consumption a more normalized activity to be done in moderation.*

+3 strongly agree ----- -3 strongly disagree

(2) *18 year olds are more susceptible to binge drinking and other risky/irresponsible behaviors than people who are 21 and older.*

-3 strongly agree ----- +3 strongly disagree
(note reversed polarity)

The basic statistical information about the two experimental groups is given in Table 1 and the tweet distribution of each participant in Group-1 is shown in Figure 1. Participants are denoted by a

two-letter abbreviation (WS, EP and so on). The current data set is only a fraction of a larger corpus, which is currently under development. Additional datasets cover a variety of discussion topics and involve different groups of participants.

Group	# participants	# tweets	Influencer
1	20	225	WS
2	14	222	EP

Table 1: Selected details of two experimental groups.




Figure 1: Tweet distribution for each participant in Group-1 where participants with asterisk are against “lowering drinking age”.

As we would like to know the participants’ pre- and post-discussion attitudes about the overall topic, we used the responses on 11 survey questions to calculate how strongly participants feel on the overall topic of discussion. Each question is given on a seven-point scale ranging from “+3” to “-3”, where “+3” implies strongly agree to keep drinking age at 21 and “-3” means strongly disagree. Positions of participants are determined by adding the scores of the 11 questions according to their responses on pre- or post- discussion questionnaires. Figure 2 is an example of pre-discussion responses for two participants in Group-1. WS largely agrees that drinking age should be kept at 21 whereas EK has an opposing opinion. The pre- and post-discussion attitudes of participants in Group-1 towards the overall topic are shown in Figure 3.

	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	Q11	Sum
WS	3	0	2	-2	0	-1	3	3	3	0	0	11
EK	-2	-2	1	2	0	-2	-3	0	0	-3	-3	-12

Figure 2: Pre-discussion survey scores of WS and EK.

Subsequently, we computed relative *pre-discussion attitude distance* between each participant and the influencer based on the pre-discussion surveys and their *post-discussion attitude distance* based on the post-discussion surveys. We normalized these distances to a $[0, 2]$ interval to be consistent with cosine distance computation scale used in the TPN module. The changes from pre-discussion attitude distance to post-discussion attitude distance based on the surveys are considered the gold standard against which the system-computed TPN values are measured. As shown in Figure 4(a), the pre-discussion distance between WS and EK is 1.43 (first bar) and the post-discussion distance is 0.07 (second bar), which implies their positions on the overall topic moved significantly closer. We also note that WS’s position did not change much throughout the discussion (Figure 3). This was just as we expected since WS was our designated influencer, and this fact was additionally confirmed in the post survey: in response to the question “Who was the influencer in the discussion?” the majority of participants selected WS. The post survey responses from the other group also confirmed our selected influencer. In addition, we used the automated DSARMD system (Strzalkowski et al., 2013) to compute the most influential participants in each group, and again the same people were identified.




Figure 3: Pre- and post-discussion attitudes of participants in Group-1 where the left bar of the participant is their pre-discussion attitude and right bar of the participant is their post-discussion attitude.

4.2 Experiment




Figure 4: (a) Relative position change between speakers WS (the influencer) and EK based on surveys and automatically computed TPN distance. The first bar in each pair corresponds to their pre-discussion distance and second bar is post-discussion distance. We note that TPN predicts correctly that WS and EK move closer together. (b) Relative position change between participants WS and BC.

After detailed analysis of participants' opinion before and after the discussion, two twitter discussions are run through our system to extract the required information in order to compute topical positioning as explained in section 3. In Group-1, ten meso-topics were generated by our system (including, e.g., “drinking age”, “teens” and “alcohol”). Each participant's polarity associated with these meso-topics was computed by our system to form ten-dimensional topical positioning vectors for Group-1. In our experiment, we used the first quarter of discussion to compute initial topical positioning of the group and last-three quarters to compute the final topical positioning. Once the pre- and post-topical positioning were determined, the topical positioning map between participants was calculated accordingly, i.e., pre- and post-TPN. We used the first quarter of discussion for the initial TPN because we required a sufficient amount of data to compute a stable measure; however, we expected it would not fully represent participants' initial positions. Nonetheless, we should still see the change when compared with post-TPN, which was computed on the last three-quarters of the discussion. In order to detect the opinion shifts and also to measure the effect of the influencer on a group, we tracked the changes in the TPN with respect to the influencer. As shown in Figure 4(a), the pre-TPN between WS and EK is 1.33 (third bar) and post-TPN is 0.72 (fourth bar). Hence, the system determines that their opinions are moving closer which conforms to the survey results. Figure 4(b) is another example of WS and BC that system result shows the same tendency as the survey result. The pre-discussion distance between WS and BC is 1.87 (first bar) and the post-discussion distance is 1.42 (second bar), which implies their positions on the overall topic moved closer after discussion. In system detection, the pre-TPN between is 1.56 (third bar) and post-TPN is 1.02 (fourth bar), which also concludes their attitudes are closer. Another examples showing that speaker moved away from influencer are in Figure 5(a) and 5(b). According to the survey, the pre-discussion attitude distance between WS and CC is 0.62 (first bar) and post-discussion attitude distance is 1.35 (second bar), which implies their positions diverged after the discussion. Our system determined pre-TPN between WS and CC is 1.0 (third bar) and post-

TPN is 1.29 (fourth bar), which shows their divergence.

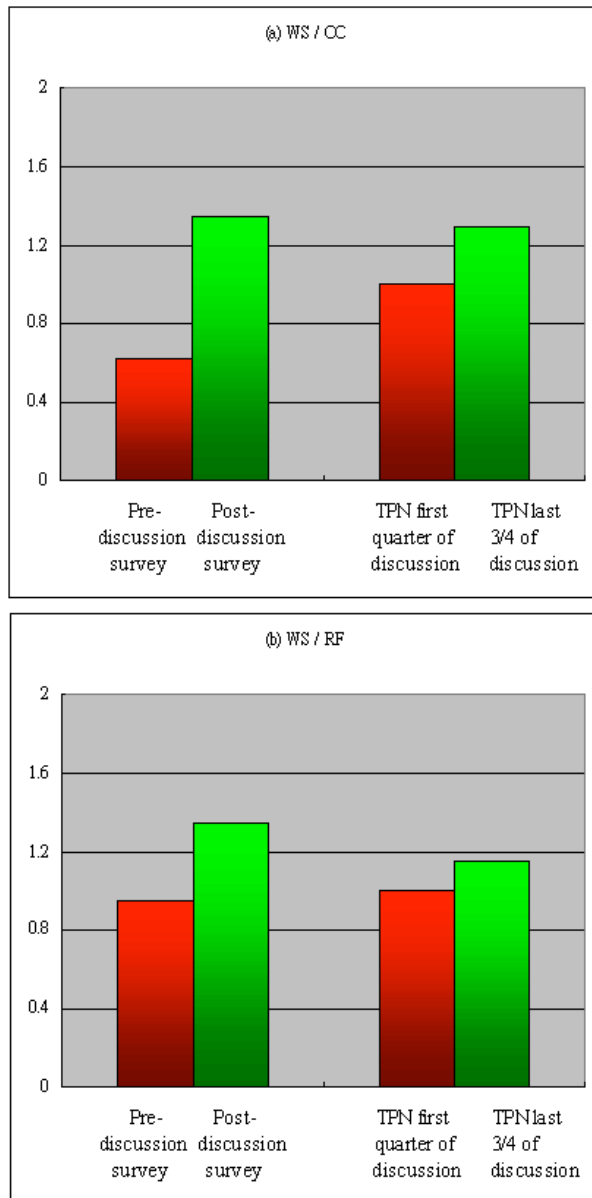


Figure 5: (a) Relative position change between WS and CC based on surveys and TPN. (b) Relative position change between participants WS and RF. We note that RF moves away from WS, which is correctly predicted by TPN.

In a separate exercise we also explored different parts of the Twitter session to compute pre-TPN and post-TPN, in addition to the $\frac{1}{4}$ vs. $\frac{3}{4}$ split discussed above. In particular, we computed TPN distances between speakers at first $\frac{1}{2}$ vs. second $\frac{1}{2}$,

first $\frac{1}{4}$ vs. last $\frac{1}{4}$, first $\frac{1}{3}$ vs. last $\frac{1}{3}$, etc. Experiment results show that using the first quarter of discussion as initial topical positioning and last quarter as final topical positioning ($\frac{1}{4}$ vs. $\frac{1}{4}$) produces the most accurate prediction of opinion changes for all group participants: 87.5% in Group-1 and 76% in Group-2. We should also note here that there is no specific correlation between the meso-topics and the overall topic other than the meso-topics arise spontaneously in the conversation. The set of meso-topics in the second discussion on the same topic was different than the in the first discussion. In particular, meso-topics are not necessarily correlated with the aspects of the overall topic that are addressed in the surveys. Nonetheless, the TPN changes appear to predict the changes in surveys in both discussions. At this time the results is indicative only. Further experiments need to be run on additional data (currently being collected) to confirm this finding.

5 Conclusion

In this paper, we described an automated approach to detect participant's Topical Positioning and capture the opinion shifts by Topical Position Maps. This work is still in progress and we intend to process more genres of data, including Twitter and on-line chat, to confirm effects seen in the data we currently have. The future work should be able to account for the relationship between meso-topic and overall topic (i.e., supporting meso-topic means for or against overall topic). A potential solution could be determined by aligning with TPN of influencers who are known strongly pro- or against- overall topic. Another avenue of future work is to apply proposed model on virtual chat-room agent to guide the discussion and change participants' attitudes.

References

- Bradley, M. M., & Lang, P. J. 1999. *Affective norms for English words (ANEW): Instruction manual and affective ratings*(Tech. Rep. No. C-1). Gainesville, FL: University of Florida, The Center for Research in Psychophysiology.
- Broadwell, George, Jennifer Stromer-Galley, Tomek Strzalkowski, Samira Shaikh, Sarah Taylor, Umit Boz, Alana Elia, Laura Jiao, Ting Liu, and Nick Webb. "Modeling Socio-Cultural Phenomena in Dis-

- course." *Journal of Natural Language Engineering* (2012): 1-45.
- David, S., and Pinch, T. J. 2006. Six degrees of reputation: The use and abuse of online review and recommendation systems. *First Monday*. Special Issue on Commercial Applications of the Internet.
- Fellbaum, C., B. Haskell, H. Langone, G. Miller, R. Poddar, R. Teng and P. Wakefield. 2006. *WordNet 2.1*.
- Hu, M., and Liu, B. 2004. Mining opinion features in customer reviews. *In Proceedings of AAI*, 755–760.
- Klein, D., & Manning, C. D. 2003. Accurate unlexicalized parsing. *In Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1* , 423-430. Association for Computational Linguistics.
- Martineau, J., & Finin, T. 2009. Delta tfidf: An improved feature space for sentiment analysis. *In Proceedings of the 3rd AAI International Conference on Weblogs and Social Media*, 258-261.
- Philip Beineke, Trevor Hastie, Christopher Manning and Shivakumar Vaithyanathan 2003. An exploration of sentiment summarization. *In Proceedings of AAI*, 12-15.
- Shaikh, Samira, et al 2012. Modeling Influence in Online Multi-party Discourse. *Cloud and Green Computing (CGC), 2012 Second International Conference on. IEEE*.
- Small, S., Strzalkowski, T. and Webb, N. 2010. *ECO: Effective Communication Online*. Technical Report ILS-015, University at Albany, SUNY
- Somasundaran, S., & Wiebe, J. 2009. Recognizing stances in online debates. *In Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*.
- Strapparava, C., and Mihalcea, R. 2008. Learning to Identify Emotions in Text. *In Proceedings of the ACM Conference on Applied Computing ACM-SAC*.
- Strzalkowski, T.; Broadwell, G. A.; Stromer-Galley, J.; Shaikh, S.; Taylor, S.; and Webb, N. 2010. Modeling socio-cultural phenomena in discourse. *In Proceedings of the 23rd International Conference on Computational Linguistics*, 1038-1046.
- Strzalkowski, T., Samira Shaikh, Ting Liu, George Aaron Broadwell, Jennifer Stromer-Galley, Sarah M. Taylor, Veena Ravishankar, Umit Boz, Xiaoi Ren: Influence and Power in Group Interactions. *SBP 2013*: 19-27
- Vogt, T., Andre', E., & Bee, N. 2008. EmoVoice—A framework for online recognition of emotions from voice. *Perception in Multimodal Dialogue Systems*, 188-199.
- Wiebe, J., Wilson, T., and Cardie, C. 2005. Annotating expressions of opinions and emotions in language. *Journal of Language Resources and Evaluation* 39(2-3):165–210
- Zhuang, L., Jing, F., Zhu, X. Y., & Zhang, L. 2006. Movie review mining and summarization. *In Conference on Information and Knowledge Management: Proceedings of the 15 th ACM international conference on Information and knowledge management*, 43-50.