Gaining and Losing Influence in Online Conversation

Arun Sharma and Tomek Strzalkowski
ILS Institute, State University of New York – University at Albany
1400 Washington Ave, Albany, NY 12222 USA
asharma5@albany.edu, tomek@albany.edu

Abstract

In this paper, we describe a study we conducted to determine, if a person who is highly influential in a discussion on a familiar topic would retain influence when moving to a topic that is less familiar or perhaps not as interesting. For this research, we collected samples of realistic on-line chat room discussions on several topics related to current issues in education, technology, arts, sports, finances, current affairs, etc. The collected data allowed us to create models for specific types of conversational behavior, such as agreement, disagreement, support, persuasion, negotiation, etc. These models were used to study influence in online discussions. It also allowed us to study how human influence works in online discussion and what affects a person's influence from one topic to another. We found that influence is impacted by topic familiarity, sometimes dramatically, and we explain how it is affected and why.

Keywords: Natural Language Processing, Computational Sociolinguistics, Social Computing, Influence Detection, Text Analysis, Data Analytics

1. Introduction

This research was undertaken to help us understand how influential people behave in group conversations when they discuss the topics about which they have less knowledge. How does the topic of conversation affect their behavior in a group, specifically in an online conversation? Since a great deal of our communications occurs online, it is important to study and understand the sociolinguistic behavior of people who have high degree of influence in such environments.

For this research, we recruited participants in groups of 3 to 5 to discuss topics that they are all familiar with. Once we identified the most influential participant in each group, we moved them to a new group of 3 to 5 participants that discussed a different topic that our influencers were less familiar with. We collected interactions from both rounds of discussions as our data set. We automatically analyzed the data for various sociolinguistic factors and computed the degree of influence for each participant, using the DSARMD toolkit (Broadwell et al., 2012). Some of the factors considered included the rate of topic introduction, participation frequency, as well as patterns of language use, dialogue acts, references to people and other named entities (Named Entity Tagging), etc. These factors were further aggregated into predictions of sociolinguistic behavior in conversation, such as Argument Diversity, Disagreement Measure, Network Centrality, Topic Control, Involvement etc.

2. Related Work

Our objective is to study how people who are influential in one context retain or lose their influence in another context. We focused on the topic of conversation, and specifically the participants' degree of interest and familiarity with it. Research on measuring influence is a relatively new area in computational linguistics and natural language processing; although it has some commonalities with automated sentiment analysis and dialogue understanding, which have much longer history. The key recent work that the present study draws on is (Strzalkowski et al., 2013) and (Broadwell et al, 2012). This research developed and validated a computational approach for modeling sociolinguistic behaviors in conversation within small groups of up to 30 participants. Furthermore (Shaikh et al. 2013) expanded this work to track influence dynamics in conversation and how people's opinions change under influence. Our current research has utilized the approach and methods discussed in these works. Specifically, we use the components of sociolinguistic toolkit to compute and compare influence of all participants in our experiments and then determine how it relates to their knowledge of the discussion topic.

3. Methodology (Hypotheses, Study Design and Procedures)

The following sections describe how the study was undertaken, online survey formed, and responses collected.

3.1 Research Hypothesis

Our research starts with the following hypothesis:

"Individuals who have a high degree of influence within a group discussing a topic they are knowledgeable about, remain influential when moved to another group that discusses a different topic they are less knowledgeable about."

If this hypothesis could be confirmed, it would suggest that personal capabilities and behavioral profile largely determine a person's influence in conversation. Conversely, if the hypothesis is rejected, it would mean that one's influence is partly conditioned upon the topic of conversation and the relative familiarity with it. While the experiment described here is on a relatively small scale, the result would have consequences on influence modeling in general, including in large scale social networks where information diffusion is often related to a person's degree of influence. (e.g., Hofman et al, 2017)

3.2 Description of the Experiment

In the first part of our study, we asked the recruited participants to answer open-ended questionnaire about possible discussion topics of their interest. The questions focused on how much they like to talk about a subject of their interest. How much knowledge they think they have in that area of their interest.

After the data analysis of the first phase we selected the respondents based on their answers on the questionnaire and placed them in groups of 3 to 5 people to discuss a selected topic that matches their interests. These online discussions took around 1 to 1.5 hours.

In the second phase, we moved participants between groups so that the most influential participant from each phase 1 group was placed in a new group that was to discuss a topic they were less familiar with. Another round of discussions of 1-1.5 hours duration was conducted in each new group.

After all data was collected, the transcript of each discussion was analyzed by the DSARMD software that assigns a degree of influence to each participant in a group (Broadwell et al., 2012). We compared the performance of each participant in phase 1 and phase 2 groups relative to their familiarity with the discussion topics. We should note that the participants generally did not know one another and were only aware of their anonymized user ids, which were assigned by the investigator. The website we used to collect the chat logs is Chatzy.com, which allows to create private chat rooms that safely store chat data. The participant chat user id, timestamp of when they post a message and the content of the message were saved as part of the logs of discussion. Participants surveys were created using available Surveymonkey.com. Survey Monkey forms were only used to create and display the questionnaire, the responses were redirected to be saved on our server. This data was then analyzed to prove or disprove the hypothesis.

3.3 Survey Formation and Topic Selection

To collect student responses a survey was formed. First topic selection was done by coming up with specific topics for the survey. For each topic, the participants rated their interest and familiarity on a 5-point Likert scale (Table 1).

The topics were selected from areas where participants (graduate students at a public university) likely had opinions and knowledge. Some of the topics selected were from Sports, Movies, Academics, Current Affairs etc.

Not at all knowledgeable at all	1
Less than an average person	2
As knowledgeable as anybody	3
Probably more than most people	4
A lot more than most people	5

Table 1: *Measuring Knowledge for a topic*

The participants who selected option 4 or 5 were then asked if they would be interested in a group discussion for each topic (Table 2).

Not interested at all	1
Not particularly interested	2
Moderately interested	3
Quite interested	4
Very much interested	5

Table 2: *Options for Interest in Group Discussion*

The participants who indicated an interest in discussion, were further asked to write in a few words what aspects of the topic they would like to talk about. They were also asked if they would be interested in leading a discussion.

The survey covered 10 topics with up to 4 questions per topic, and it generally took approximately 15 minutes to complete.

3.4 Surveys and Group Formation

This study involved 54 participants, all graduate students at University at Albany. Data collected from the surveys were collated per topic and per participant in order to (1) identify topics that were of interest to at least some participants; and (2) form phase 1 groups based on topic interest. We required at least 3 participants per group.

After analysis of the Survey data collected from participant's 10 discussion groups were formed.

The conversation data from the discussions were collected in the text format for further analysis and to compute influence scores for participants, and rank them by the degree of influence. The most influential participants in each group were re-assigned to participate in another online group discussion this time on a topic they were less knowledgeable about. They were placed with people who had indicated they were knowledgeable about this new topic and were interested to discuss it. Again, the conversation was recorded, influence scores were computed for each participant, and participants were ranked by their degree of influence.

3.5 Demographic Makeup

Taking into consideration the privacy, limited personal data was collected from individual participants of the study. The overall demographic makeup of 54 participants was 70.37% of the participants were males and 29.63% participants were female.

4. Analysis of Online Discussion Data

The online conversation was analyzed automatically by a suite of tools developed under the DSARMD Project (Broadwell et al, 2012).

4.1 Finding Influencers based on predictions of Social Behavior in Conversation

The participants were ranked based on their influence, thus helping us select the influencer in each group along with scoring of every participant on various component metrics such as Network Centrality, Disagreement Measure, Topic Control, Involvement, among others, that contribute to the assessment of the degree of influence (Strzalkowski et al., 2013). All these attributes were computed for each conversation and participant. We describe the key metrics below. For a more detailed explanation, and how the scores are combined, the reader is referred to (Broadwell et al., 2012) and (Strzalkowski et al., 2012).

Network Centrality (NC): It is the measure of degree to which a participant in the conversation is a "center" of communication in that group. A participant has a high degree of Network Centrality when other participants address more of their utterances towards her or him, and whose topics are most discussed by others.

Disagreement Measure (DM): Disagreements with others is a way to control the topic of conversation by way of identifying or correcting what the participant sees as a problem. The more disagreement a participants shows, relative to other participants, the higher his/her Disagreement Measure.

Topic Control (TC): Topic Control reflects a conversational behavior where participants attempt to impose a topic of conversation. This may be accomplished by introduction of preferred topics that are subsequently discussed at length by others, or by successfully continuing discussion of selected topics. The ability to introduce topics and make others talk about them indicate the degree of topic control in the conversation.

Involvement (Inv): This behavior reflects the degree of engagement in a discussion measured as the proportion of conversational turns contributed by each participant.

Argument Diversity (AD): Participants who use a broader range of arguments in the conversation has a higher degree of Argument diversity. This measure includes amongst others, the size of one's vocabulary, usage of specialized terminology, and citing authoritative sources.

4.2 Results of Phase 1 experiments

Group discussions were conducted where an influential participant was identified in every group. We note again that in Phase 1 all groups were composed of participants who identified themselves as knowledgeable about the discussion topic. Below we show details of Influence (Inf) and the component metrics for two Phase 1 groups. Other groups had similar distributions of scores.

Group 1: U.S. Immigration System and Reforms
Number of Participants: 4

Table 3 shows performance of participants in Group 1, with Person 3 identified as the Influencer.

User	Inf	NC	DM	TC	AD	Inv
Person1	0.41	0.103	0.397	0.125	0.291	0.410
Person2	0.19	0.226	0.170	0.125	0.275	0.334
Person3	**0.70**	**0.219**	**0.379**	**0.223**	**0.307**	**0.347**
Person4	0.08	0.138	0.023	0.055	0.103	0.170

Table 3: *Participant's performance in Group 1 discussion. Inf is overall influence score; the other columns show scores for component measures: Network Centrality (NC), Disagreement Measure (DM), Topic Control (TC), Argument Diversity (AD), and Involvement (Inv)*

Group 2: Globalization
Number of Participants: 5

Table 4 shows performance of participants in Group 2, with Person 2 and Person 5 both identified as influencers. Note that their Inf scores are very close.

User	Inf	NC	DM	TC	AD	Inv
Person1	0.15	0.232	0.095	0.127	0.193	0.195
Person2	0.45	0.346	0.257	0.288	0.318	0.376
Person3	0.04	0.0	0.095	0.002	0.104	0.196
Person4	0.33	0.185	0.215	0.101	0.237	0.319
Person5	0.46	0.235	0.334	0.085	0.138	0.178

Table 4: *Participants' performance in Group 2 discussion*

Similarly, the influencers were identified from the remaining online group discussions. In most cases, a single person was identified as the influencer. In two cases, no one was selected (all scores were closely distributed around the mean).

4.3 Results of Phase 2 experiments

In the second phase of the study, the persons identified as influential in each Phase 1 group was placed in a new group where the topic of discussion was less familiar to them. To be clear: the new group was composed of participants who were knowledgeable about the discussion topic, although not particularly influential (based on their Phase 1 performance); however, our Phase 1 influencer now assigned to this group was less familiar with the topic, according to the participant's self-assessment in the survey. We also had to discard two Phase 1 group discussion datasets due to relatively low level of participation by the members of these groups, and thus system's inability to identify influencers.

The influence scores were computed again along with their components (see the next section). These influence scores should be analyzed along with all the component scores, in order to see the primary source of each person influence or lack thereof.

4.4 Outcomes

The following results were obtained after the second round of discussions. Every influencer was evaluated using the same measures and the results were studied. Below we show how selected Phase 1 Influencers (1 through 6) fared in Phase 2 discussions.

Influencer 1
Knowledgeable about: U.S. Immigration Reform
Less Knowledgeable about: Science-Fiction Movies

Table 3 shows performance of Influencer 1 (Person 3) in the first-round discussion discussing US Immigration Reform. We note that Influencer 1 had high scores on all component measures, and the highest scores in topic control and argument diversity.

User	Inf	NC	DM	TC	AD	Inv
Person1	0.42	0.140	0.458	0.104	0.218	0.252
Person2	0.16	0.266	0.091	0.108	0.217	0.313
Person3	0.13	0.066	0.036	0.078	0.218	0.261
Person4	0.69	0.377	0.402	0.276	0.318	0.349

Table 5: *Performance of Influencer 1 (Person 3) in the second-round group discussing science-fiction movies.*

Table 5 shows Influencer 1's performance in the second-round discussion. This time, our influencer (Person 3 again) was amongst the least influential people in this group with the lowest network centrality score, and the lowest topic control score – both related to the topic familiarity. Note that the involvement measure and argument diversity, which are more related to behavioral attributes of a person, remain relatively high, but cannot compensate for the loss of topical scores.

Influencer 2
Knowledgeable about: Globalization
Less Knowledgeable about: Science-Fiction Movies

This participant was given the ID "Person 5" in first round -group (see Table 4) and "Person 3" in the second-round group (Table 6). As an influencer in the first-round discussion, Person 5 had highest disagreement measure (DM), the second highest score in network centrality (NC), a moderate amount of topic control, and argument diversity and was less involved.

User	Inf	NC	DM	TC	AD	Inv
Person 1	0.528	0.087	0.295	0.170	0.344	0.396
Person 2	0.440	0.305	0.336	0.165	0.259	0.363
Person 3	0.268	0.098	0.285	0.083	0.166	0.208
Person 4	0.118	0.089	0.027	0.059	0.197	0.262

Table 6: *Performance of Influencer 2 (Person 3) in the second-round group discussing science-fiction movies.*

Influencer 2 was not quite influential in the second-round group. While the loss of influence was not as dramatic as for Influencer 1, the main reason here was the loss of Network Centrality status. Note that Influencer 2, unlike Influencer 1, was not a strong driver of conversation topic (Topic Control measure); instead his influence derived mostly from other people deferring to his superior topic expertise, which of course was not present in round 2.

Influencer 3
Knowledgeable about: Globalization
Less Knowledgeable about: Science Fiction Movies

This participant was given the ID "Person 3" in the first-round group (Table 7) and "Person 2" in the second-round group (Table 8)

User	Inf	NC	DM	TC	AD	Inv
Person 1	0.204	0.355	0.060	0.135	0.259	0.246
Person 2	0.272	0.235	0.296	0.200	0.402	0.507
Person 3	0.860	0.094	0.592	0.120	0.249	0.306

Table 7: *Performance of Influencer 3 (Person 3) in the first-round group discussing globalization.*

Influencer 3 had the highest disagreement measure (DM), was second in involvement and low in all other measures.

User	Inf	NC	DM	TC	AD	Inv
Person 1	0.539	0.321	0.321	0.316	0.504	0.457
Person 2	**0.321**	**0.199**	**0.321**	**0.222**	**0.225**	**0.311**
Person 3	0.596	0.478	0.357	0.061	0.258	0.364

Table 8: *Performance of Influencer 3 (Person 2) in the second-round group discussing science-fiction movies. Note that Person 1 and Person 2 were the most influential participants in this group as they were very close to each other in influence score.*

In the second round discussion, Influencer 3 (as Person 2) was the least influential in this group. While he ranked 2nd in disagreement measure (DM) and topic control (TC), his network centrality score (NC) was the lowest and so was Involvement (Inv). In this case, the loss of influence is attributed to the relative decline of disagreement and network centrality measures.

Influencer 4
Knowledgeable about: Criminal justice system in the U.S.
Less Knowledgeable about: Science Fiction Movies

This participant was given the ID "Person 1" in the first-round group (Table 9) and "Person 2" in the second round group (Table 10)

User	Inf	NC	DM	TC	AD	Inv
Person 1	**0.819**	**0.211**	**0.465**	**0.238**	**0.433**	**0.461**
Person 2	0.233	0.219	0.153	0.117	0.181	0.181
Person 3	0.118	0.171	0.037	0.079	0.158	0.195
Person 4	0.275	0.321	0.089	0.160	0.216	0.312

Table 9: *Performance of Influencer 4 (Person 1) in the first-round group discussing criminal justice system.*

Influencer 4 had highest disagreement measure (DM) in the first round group. He also had highest topic control score (TC), highest involvement score (Inv) and argument diversity (AD). Person 1 was clearly the most dominant in this group.

In the second round discussion, Influencer 4 was only the 3rd influential person in the group. He had lowest score on the disagreement measure (DM), 2nd in network centrality scores (NC), while still a high topic control (TC) and highest involvement (Inv) and argument diversity (AD) in the group.

User	Inf	NC	DM	TC	AD	Inv
Person 1	0.250	0.360	0.102	0.232	0.294	0.321
Person 2	**0.257**	**0.329**	**0.097**	**0.221**	**0.327**	**0.371**
Person 3	0.274	0.051	0.198	0.053	0.171	0.215
Person 4	0.645	0.196	0.595	0.067	0.163	0.264

Table 10: *Performance of Influencer 4 (Person 2) in the first-round group discussing science fiction movies.*

Nonetheless, the group was now dominated by Person 4 who, with the very high disagreement score clearly managed to block most of Person 2's attempts at topic control, in spite of his relatively high network centrality.

Influencer 5
Knowledgeable about: Science Fiction Movies
Less Knowledgeable about: U.S. Immigration Reform
This participant was given the ID "Person 1" in the first and also in the second round group (Tables 11 and 12).

User	Inf	NC	DM	TC	AD	Inv
Person 1	**0.54**	**0.321**	**0.321**	**0.316**	**0.504**	**0.457**
Person 2	0.32	0.199	0.321	0.222	0.225	0.311
Person 3	0.59	0.478	0.357	0.061	0.258	0.364

Table 11: *Performance of Influencer 5 (Person 1) in the second-round group discussing science-fiction movies.*

Influencer 5 was one of the two influential participants in the group (Person 3 was the other). He had highest argument diversity (AD), topic control (TC), and involvement (Inv) scores, while the other influencer had top scores in network centrality (NC) and disagreement (DM). This group is somewhat marginal in the sense that no single participant dominated the conversation.

User	Inf	NC	DM	TC	AD	Inv
Person1	**0.290**	**0.402**	**0.179**	**0.245**	**0.362**	**0.414**
Person2	0.033	0.0	0.0	0.0	0.074	0.066
Person3	0.662	0.124	0.568	0.122	0.316	0.335
Person4	0.412	0.273	0.226	0.170	0.203	0.069

Table 12: *Performance of Influencer 5 (Person 1) in the second-round group discussing US Immigration Reform.*

In the second-round group, Influencer 5 was a distant third in influence ranking. While he retained high scores in network centrality (NC) and involvement (Inv), other metrics have fallen significantly, esp. disagreement (DM).

Influencer 6
More Knowledgeable about: Science Fiction Movies
Less Knowledgeable about: U.S. Immigration Reform

This participant was given the ID "Person 4" in the first-round group and "Person 3" in the second-round group.

User	Inf	NC	DM	TC	AD	Inv
Person1	0.42	0.140	0.458	0.104	0.218	0.252
Person2	0.16	0.266	0.091	0.108	0.217	0.313
Person3	0.13	0.066	0.036	0.078	0.218	0.261
Person4	**0.69**	**0.377**	**0.402**	**0.276**	**0.318**	**0.349**

Table 13: *Performance of Influencer 6 (Person 4) in the first-round group discussing science-fiction movies.*

Influencer 6 had not only the highest influence score in this group, but he had the highest scores on all component measures. Clearly, this person was extremely good at discussing science fiction movies. And yet, this dominating influence all but vanished when he was part of the second-round group discussing immigration matters (Table 14).

User	Inf	NC	DM	TC	AD	Inv
Person1	0.34	0.425	0.361	0.343	0.414	0.435
Person2	0.17	0.096	0.180	0.104	0.28	0.339
Person3	**0.04**	**0.035**	**0.011**	**0.013**	**0.070**	**0.335**
Person4	0.11	0.067	0.361	0.044	0.204	0.069

Table 14: *Performance of Influencer 6 (Person 3) in the second-round group discussing U.S. Immigration Reform.*

In the second-round discussion, Influencer 6 had the lowest influence score, and moreover he had the lowest scores on all component measures. This is a remarkable case of a complete loss of influence, which is clearly not due to the lack of participation: we note that Person 3 remains highly involved in discussion (high Inv score), but her impact is marginal.

In all cases covered by this study, a significant loss of influence was noted between round 1 and round 2. We discuss our observations next.

5. Conclusion

Based on the analysis of data that was collected during this study, we can make the following observation: People who are influential in group discussions on familiar subjects, *lose their influence* when placed in a group discussing a less familiar topic. All the influential participants saw their influence decline, often so dramatically that they become the least influential participants in a group.

Interestingly, the Involvement component score of the Influencer score matrix remains approximately the same for 4 out of 6 the top influential people from round 1 to round 2. This suggests that the decline in influence is not simply explained by the lack of participation. The Topic Control score declines in 5 out of 6 instances, which is a significant factor in the loss of influence. The Disagreement Measure has also decreased, which is another key factor. The other measures, Network Centrality and Argument Diversity showed mixed results, increasing in some cases and decreasing in others. This leads us to conclude that the loss of influence associated with topic change is largely due to loss of Topic Control, and the participants' inability to take strong positions that others may endorse even when they initially disagree with them.

Thus, based on our analyses of the group discussions we further conclude that the hypothesis **"**Individuals who have a high degree of influence within a group discussing a topic they are knowledgeable about, remain influential when moved to another group that discusses a different topic they are less knowledgeable about" is not supported by the data that we have. On the contrary, we have shown that when influential people are moved to groups where unfamiliar topics are discussed, their influence declines, sometimes significantly. This loss is attributed to the decreased scores of Topic Control and Disagreement measures, which are directly associated with the topical knowledge; however, this loss is not compensated by an increase on other measures such as Involvement and Network Centrality.

More extensive study and experiments are required to study this phenomenon more conclusively and an effort must be made to include a more varied group of participants to made data more representative and study the influence in online conversations. Also, other domains like online networking websites and discussion forums should be looked at more closely and analyzed in future studies to obtain a more representative set of data for studying influence in online world.

6. References

Shaikh, S., Strzalkowski, T., Stromer-Galley, J., Broadwell, G. A., Taylor, S., Liu, T., Boz, U. (2012). Modeling Influence in Online Multi-Party Discourse. 2012 Second International Conference on Cloud and Green Computing. doi:10.1109/cgc.2012.94

Strzalkowski, T., Shaikh, S., Liu, T., Broadwell, G.A., Stromer-Galley, J., Taylor, S. M., Ravishankar, V., Boz, U., Ren, X.: Influence and Power in Group Interactions. SBP 2013: 19-27

Broadwell, G. A., Stromer-Galley, J., Strzalkowski, T., Shaikh, S., Taylor, S., Liu, T., Webb, N. (2012). Modeling Sociocultural phenomena in discourse. Natural Language Engineering,19(02),213-257. doi:10.1017/s1351324911000386

Shaikh, S., Strzalkowski, T., Taylor, S., Lien, J., Liu, T., Broadwell, G. A., Peshkova, Y. (2015). Understanding Cultural Conflicts using Metaphors and Sociolinguistic Measures of Influence. Proceedings of the Third Workshop on Metaphor in NLP. doi:10.3115/v1/w15-1408

Strzalkowski, T., Shaikh, S., Ravishankar, V., Stromer-Galley, J., Crowley. J, Lin, C.S. (2013). Topical Positioning: A New Method for Predicting Opinion Changes in Conversation. Proceedings of the Workshop on language in Social Media (LSAM 2013), pages 41-48.

Strzalkowski, T., Broadwell, G.A., Stromer-Galley, J., Shaikh, S., Taylor, S., Webb, N.: Modeling socio-cultural phenomena in discourse. In: Proceedings of the 23rd International. Conference on Computational Linguistics, Beijing, China (2010)

Strzalkowski, T., Broadwell, G.A., Stromer-Galley, J., Shaikh, S., and Taylor, S. (2012) Modeling Leadership and Influence in Online Multi-Party Discourse. COLING-2012 Conference, Mumbai, India.

Hofman, J.M., Sharma, A., and Watts, D.J.. (2017). Prediction and explanation in social systems. Sciencemag.org. *Science* 03 Feb 2017: Vol. 355, Issue 6324, pp. 486-488 DOI: 10.1126/science.aal3856