In You We Follow: Determining the Group Leader in Dialogue

David B. Bracewell and Marc T. Tomlinson

Language Computer Corporation
Richardson, TX
{david,marc}@languagecomputer.com

Abstract. In this paper, we investigate whether the social roles of dialogue participants can be recognized through the social actions performed by the participant in their interactions with others in the group. Specifically we focus on determining if a participant is the leader of the group. We decompose the problem into identifying the social goals for participant discourse segments. These social goals are represented through a set of eleven psychologically-motivated social acts. We then model leadership using a sociological-inspired model called social rank which takes into account the social capital accumulated by the participant over the course of a single dialogue. We explore these models in task-oriented dialogues communicated in English, Arabic, and Chinese and show that the incorporation of social rank can improve precision of detecting the leader by 14% in English, 8% in Arabic, and 4% in Chinese.

1 Introduction

Leaders in task-oriented groups provide guidance, push for success, and facilitate discussion. The group affords a leader power granting the leader influence and status within and over the group. Identification of group leaders and leadership qualities is an important first-step in understanding the power structure of a group.

Much attention has been given to pure network-based approaches (e.g. PageRank [1]) to determining the power and influence an individual wields within an online community [2, 3]. Recently, work in natural language processing has begun addressing the question of power in dialogue through analysis of the messages exchanged between individuals [4–7].

In this paper we examine a method for determining who the leader of a group is using the interactions contained within a single dialogue. We determine the leader through examination of the social actions performed by the dialogue participants as evidenced through the linguistic expressions they employ. Using these social actions, we rank participants using a metric we call *Social Rank* which is roughly based on the sociological theories of social capital and group stratification. We define *Social Rank* as the honor and prestige afforded to an individual by others in the group which increases their status and power within and over the group. We explore how *Social Rank* can improve over baseline and

machine learning approaches for determining the leader in dialogues communicated in English, Arabic, and Chinese.

2 Sociological Roots of Social Capital and Social Rank

The social rank of an individual is the honor and prestige afforded them by others in the group. We base our idea of social rank on the theories of group stratification formulated by Max Weber [8]. Weber's Three-component theory of stratification examines how wealth, power, and prestige work together in determining the social power of an individual [8]. Weber defines wealth as the economic resources available to an individual. Prestige is defined as the respect a person is afforded by their status or position. The final aspect of Weber's theory is power. Power can come via a variety of ways. Realizations of power come in two main types, formal power (power given to an individual by an authority) and informal power (power based on an individual's characteristics, e.g. expertise, skills; [9]). In the case of identifying the leader, we are looking for power over the task and discussion of the group.

One way in which social wealth and prestige can be determined is through an individual's social capital. Social capital originates from the field of Sociology. A variety of definitions for social capital can be found in literature. The definition of particular interest to task-oriented discourse and leadership revolves around the role that cooperation and trust play in the collective results of the group. We base our definition of social capital on Coleman [10], who describes social capital as an employable resource of individuals. He states that social capital is gained through the changing relations of individuals and in particular those changes that cause action. Coleman describes three forms of social capital: (1) *Obligation*, (2) *Information*, and (3) *Norms and Effective Sanctions*.

Obligation based social capital is gained by performing an action for another where there is an expectation and trust that the other person will repay the action. An individual is able to collect on the social capital they gained at some point in the future, i.e. "call in a favor." Information based social capital arises when one uses their network to gain information, or expert knowledge, about a topic. For example, someone interested in technology, but who does not have the time to keep up with the latest trends, might use their friends to gain information. Unlike the obligation form of social capital, information-based social capital has no expectation of being repaid in the future. Coleman's final form of social capital relates to the social norms and effective sanctions that govern an individual's actions. One of the more powerful forms of norms arises when individuals act in the interest of the group instead of their own self-interest. The rewards for adhering to this norm are often increased status, honor, and social support. In the case of task-oriented groups this may be an increase in power or even leadership.

3 Computational Model of Social Capital and Social Rank for Dialogue

An individual's desire to reach their goals drives their predisposition to belong to a group. A negotiation is performed between group members over the social identity they assume within the group, which in the process transforms individual collaborators into a collective and cooperative group [11]. Based on these social identities, individuals communicate and make social contacts within their negotiated role. The social contacts affect individual and group productivity [12]. Furthermore, individuals build social capital based on the amount and type of contact they make whilst in the group.

3.1 Social Actions

While the social contacts made by individuals in social networks are often explicit through friending, liking, disliking, etc., in dialogue they are more likely to be manifested through linguistic expressions of social intentions. We label the social intentions of an utterance using social acts. Social acts are pragmatic speech acts that signal a dialogue participant's social intentions.

For calculation of social capital and social rank we have defined three categories of social acts: *Cooperation*, *Support*, and *Hostility*. Cooperation and support are designed to roughly capture the three forms of social capital discussed in section 2. Cooperation is captured through the social acts of agreement, offering gratitude, mediation, and solidarity. Support is captured through agreement, group affordance, solidarity, and supportive behavior [13]. Finally, hostility is used to capture the lack of prestige, social wealth, and power and individual has within the group. Hostility is captured through social actions of undermining, disrespect, relationship conflict, task conflict, and challenges of credibility. We use the same definitions of and methodology for identification of these social acts in text as [14, 15].

3.2 Social Capital

Based on Coleman's [10] work, we have defined a measure of social capital based on the actions performed by individuals in a single dialogue. In particular, we posit that individuals who have higher levels of interactivity, cooperation, and support will have higher amounts of social capital. Furthermore, an individual's social capital is increased by the capital they can collect from others due to obligations, support, etc. More formally, we calculate social capital ($SCap$) for an individual P_i as:

$$SCap(P_i) = \alpha \cdot I(P_i) + \beta \cdot C(P_i) + \frac{\delta \cdot \sum_{P_j \in A(P_i)} (S(P_j, P_i) \cdot SCap(P_j))}{N - 1} \quad (1)$$

where N is the total number of dialogue participants; $I(P_i)$ is the level of interactivity participant P_i had with the group; $C(P_i)$ is the amount of cooperation

P_i had toward the group; $S(P_J, P_i)$ is the amount of support participant P_j showed toward participant P_i; and $A(P_i)$ is the set of participants who made more affordances than detractions to P_i. The α, β, and δ parameters control the effect of the individual components (interactivity, cooperation, and support) on the social capital.

Increased interaction is one mechanism to build prestige and is needed to gain obligation and information based social capital. Moreover, we posit that a leader should have more interaction with group as they need to control the group toward their outcome. $I(P_i)$ is the interactivity of P_i which captures the breadth of a participant's interaction with the group and is calculated as:

$$I(P_i) = \frac{\sum_{P_j \neq P_i} Reply(P_i, P_j)}{N} \tag{2}$$

where $Reply(P_i, P_j)$ means that there exists a turn (t) in which participant P_i was the speaker and P_i was directing their message toward P_J (the target).

While groups require a small level of conflict to be there most productive [16], it is generally agreed that cooperative groups are more likely to reach their goals than non-cooperative ones [17]. As such, we posit that a leader is someone who shows more signs of being cooperative than not. $C(P_i)$ is the amount of cooperation P_i shows towards other group members and is calculated as:

$$C(P_i) = \frac{\sum_{t=1}^{T}(Speaker(P_i, t) \cdot C_{action}(t)}{T_{P_i}} \tag{3}$$

where $Speaker(P_i, t)$ returns 1 if P_i is the speaker of turn t, $C_{action}(t)$ is the number of social acts indicative of cooperation at turn t, T is the total number of turns in the dialogue, and T_{p_i} is the total number of turns by P_i the dialogue.

The affordance of power, or support of power, by group members to an individual is a sign of that individual's power. Moreover, with these affordances comes social capital with which an individual can borrow against. We capture these affordances through signs of support from a participant P_j to P_i for participants in the set $A(P_i)$. A participant P_j belongs to set $A(P_i)$ when the total number of times P_j employs a positive (cooperative or supportive) social act is more than they employ a negative (hostile) social act toward P_i, where cooperative, supportive, and hostile are defined in section 3. In other words, the participants in set A are those for whom P_i can gain social capital.

$S(P_j, P_i)$ is the amount of support a participant P_j (in set $A(P_I)$) shows towrd P_i and is calculated as:

$$S(P_j, P_i) = \frac{\sum_{t=1}^{T}(Speaker(P_j, t) \cdot Target(P_i, t) \cdot S_{action}(t))}{T_{P_j}} \tag{4}$$

where $Target(P_i, t)$ returns 1 if P_i is the target (i.e. who the utterance was directed toward) for turn t and $S_{action}(t)$ is the number of social acts indicative of support at turn t.

Calculating social capital is done using an iterative algorithm, similar to that used for PageRank [1] which is shown in figure 1.

Iterative Method for Calculating Social Capital

1: Initialize $SCap$ for each Person P_i at turn $t = 0$ to $\alpha \cdot I(P_i) + \beta \cdot C(P_i)$
2: $k = 0$
3: *do*
4: $\quad k = k + 1$
5: $\quad SCap(P_i; t = k) = \alpha \cdot I(P_i) + \beta \cdot C(P_i) + \frac{\delta \cdot \sum_{P_j \in A(P_i)} (S(P_j, P_j) \cdot SCap(P_i; t=k))}{N-1}$
6: *until* $\left(\sum_{P_i} SCap(P_i; t = k) - \sum_{P_i} (P_i; t = k - 1) \right) < \epsilon$

Fig. 1. The iterative method for calculating social capital

The calculation of social capital begins with each participant P_i in the dialogue having an initial value equal to their interactivity and cooperation with the group. The algorithm then performs a number of iterations updating the social capital value for each P_i based on the social capital of the other participants in set $A(P_i)$. The process completes when the total social capital for the group at iteration k changes less than ϵ from the total social capital at iteration $k - 1$.

3.3 Social Rank

Joining social capital with the concepts of Weber [8], we define social rank as a metric for determining power, prestige, and status within in a group. Within the confines of virtual online groups (excluding those taking part in a game, such as World of Warcraft), economic resources, as are traditionally discussed within the context of social rank and social class, are often of little consequence. Instead, it is the social capital of an individual that defines their wealth. An individual's social capital facilitates their ability to affect change or action that makes them wealthy in online settings. Thus, we formally define social rank ($SRank$) for a participant P_i as:

$$SRank(P_i) = \lambda \cdot SCap(P_i) - \frac{\gamma \cdot \sum_{P_j \notin A(P_i)} (H(P_j, P_i) \cdot SCap(P_j))}{N - 1} \tag{5}$$

We define the negative argument of the social rank calculation as the social detraction for participant P_i. We hypothesize that an individual who does not follow their negotiated social identity creates conflict and hostility within in the group. λ and γ are free parameters, which adjust the amount that the individual's social capital and social detraction play in determining their social rank.

Social detraction measures the amount of an individual's disregard or deviance from their agreed upon social identity throw the reactions of others in the group. The main component of the social detraction is $H(P_j, P_i)$ which is a measure of hostility, or conflict, toward P_i as exhibited by another participant P_j and is calculated as:

$$H(P_j, P_i) = \frac{\sum_{t=1}^{T} (Speaker(P_j, t) \cdot Target(P_i, t) \cdot H_{action}(t))}{T_{P_j}} \tag{6}$$

where $H_{action}(t)$ is the number of social acts indicative of hostility at turn t. The group of participants who contribute to the social detraction of P_i are those who are not in the set $A(P_i)$, i.e. those participants whose showed more negative actions than positive actions toward P_i.

4 Experimental Results

For evaluation, we gathered dialogues communicate in each language from Wikipedia discussions, web forums, blog comments, and chat transcripts. The number of dialogues used for training and testing are shown in table 1. The average number of participants across the entire dataset was 7.8 for English, 16.7 for Arabic, and 8.3 for Chinese with an average of 40.8 turns in English, 29.0 turns in Arabic, and 23.6 turns in Chinese. The training data consisted of dialogues in which a leader may not have been present. The testing data consisted only of dialogues in which there was a leader present.

Table 1. Number of dialogues used for training/development and used for testing per language.

Language	# Training Dialogues	# Testing Dialogues
English	83	75
Arabic	135	25
Chinese	353	75

We set the parameters to equally weight each factor in the social capital (as $\alpha = 0.33$, $\beta = 0.33$, and $\delta = 0.33$) and social rank ($\lambda = 0.5$ and $\gamma = 0.5$) equations. We believe that these default set of parameters should provide decent performance on a wider range of data genres. However, in the future we will explore tuning the parameters to different genres (chat, forum, blog, etc.) of data using a development set and grid search.

We compared social rank against a random baseline model (randomly picking one participant as the leader) and a motif model which discovers patterns of social acts that are indicative of leadership. The motif model was used in [4] to capture pursuits of power in dialogue. The motif model determines whether or not a participant is exhibiting leadership qualities. The participant which the motif model has the highest confidence of exhibiting leadership qualities was then determined to be the leader of the dialogue. A motif model was constructed for each of the languages using a set of Yes/No annotations over the training data which was multiply annotated. The training data consisted of a total of 425 participants in English dialogues, 1,387 participants in Arabic, and 3,136 participants in Chinese with Yes/No annotations. The inter-annotator agreement rates were 72.8% for English, 83.4% for Arabic, and 88.0% for Chinese. The testing data consisted of an average of 8.3 participants in English, 4.2 participants in

Arabic, and 11.1 participants in Chinese (these averages differ than those found in the entire dataset).

We also examined the performance when incorporating the confidence score of the motif model into the social rank calculation. The equation for social rank ($SRank$) after incorporating the motif model is as follows:

$$SRank(P_i) = TL(P_i) \cdot \left(\lambda \cdot SCap(P_i) - \frac{\gamma \cdot \sum_{P_j \in D(P_i)} (H(P_j, P_i) \cdot SCap(P_j))}{N-1} \right) \quad (7)$$

where $TL(P_i)$ is the confidence that participant P_i exhibits leadership using the motif model. Table 2 lists the accuracy for determining the **one** leader in the dialogue for the random baseline, the motif model, social rank, and the combination of social rank and the motif model.

Table 2. Comparison of accuracy results in determining the leader of a dialogue between using and not using social rank

	Random	Motif	Social Rank	Motif + Social Rank
English	12%	26%	38%	40%
Arabic	24%	48%	52%	56%
Chinese	9%	48%	48%	52%

As illustrated in Table 2, social rank increased the accuracy in determining the leader by 26% over baseline and 12% over the motif model for English, 28% over baseline and 4% over the motif model for Arabic, and 39% over baseline for Chinese. Chinese was the only language which social rank did not improve the accuracy over the motif model. This most likely due to the large size of training data used that existed to train the motif model which resulted in the model generating better confidence scores. The incorporation of the confidence that a participant was exhibiting leadership qualities generated by the motif model into social rank resulted in improved accuracy for all three languages. English improved by 2%, Arabic by 4%, and Chinese by 4% of social rank alone.

5 Conclusion

In this paper we presented a computational model for determining the group leader in a dialogue using the sociological theories of social capital and social rank. We evaluated social rank on dialogues communicated in English, Arabic, and Chinese. We showed that the social rank can drastically increase performance over baseline (up to 39% for Chinese) and provides better accuracies than a machine learning classifier. The incorporation of the classifier into the social rank computation further boosted accuracies by 2% to 4%.

Acknowledgments

This research was funded by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), and through the U.S. Army Research Lab. All statements of fact, opinion or conclusions contained herein are those of the authors and should not be construed as representing the official views or policies of IARPA, the ODNI or the U.S. Government.

References

1. Brin, S., Page, L.: The anatomy of a large-scale hypertextual Web search engine. Computer Networks and ISDN Systems **30**(1-7) (1998) 107–117
2. Brendel, R., Krawczyk, H.: Detection of Roles of Actors in Social Networks Using the Properties of Actors' Neighborhood Structure. In: 2008 Third International Conference on Dependability of Computer Systems DepCoS-RELCOMEX, IEEE (2008) 163–170
3. Fisher, D., Smith, M., Welser, H.T.: You Are Who You Talk To: Detecting Roles in Usenet Newsgroups. Proceedings of the 39th Annual Hawaii International Conference on System Sciences HICSS06 **00**(C) (2006) 59b–59b
4. Bracewell, D., Tomlinson, M., Wang, H.: A motif approach for identifying pursuits of power in social discourse. In: Semantic Computing (ICSC), 2012 IEEE Sixth International Conference on, IEEE (2012) 1–8
5. Bramsen, P., Escobar-Molano, M., Patel, A., Alonso, R.: Extracting social power relationships from natural language. In: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1. HLT '11, Stroudsburg, PA, USA, Association for Computational Linguistics (2011) 773–782
6. Danescu-Niculescu-Mizil, C., Lee, L., Pang, B., Kleinberg, J.: 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 (2012) 699–708
7. Tomlinson, M., Bracewell, D.B., Draper, M., Almissour, Z., Shi, Y., Bensley, J.: Pursuing power in arabic on-line discussion forums. In: Proceedings of the Eighth Conference on International Language Resources and Evaluation. (2012)
8. Weber, M.: The theory of social and economic organization. W. Hodge, London (1947)
9. French, J.R.P., Raven, B.: The Bases of Social Power. In Cartwright, D., ed.: Studies in Social Power. Volume 35 of Studies in social power. Institute for social research (1959) 150–167
10. Coleman, J.S.: Social capital in the creation of human capital. American journal of sociology **94**(S1) (1988) 95
11. Swann Jr., W.B., Johnson, R.E., Bosson, J.K.: Identity negotiation at work. Research in organizational behavior **29** (2009) 81–109
12. Putnam, R.: Bowling Alone: The Collapse and Revival of American Community. Touchstone Books by Simon & Schuster (2000)
13. Bracewell, D.B., Tomlinson, M., Shi, Y., Bensley, J., Draper, M.: Who's playing well with others: Determining collegiality in text. In: Proceedings of the 5th IEEE International Conference on Semantic Computing (ICSC 2011), Palo Alto, IEEE Computer Society (2011)

14. Bracewell, D.B., Tomlinson, M.T., Brunson, M., Plymale, J., Bracewell, J., Boerger, D.: Annotation of Adversarial and Collegial Social Actions in Discourse. In: 6th Linguistic Annotation Workshop. Number July (2012) 184–192
15. Bracewell, D., Tomlinson, M., Wang, H.: Identification of social acts in dialogue. In: 24th International Conference on Computational Linguistics (COLING). (2012)
16. Jehn, K.A., Mannix, E.A.: The Dynamic Nature of Conflict: A Longitudinal Study of Intragroup Conflict and Group Performance. Academy of Management Journal **44**(2) (2001) 238
17. Little, J.W.: Norms of Collegiality and Experimentation: Workplace Conditions of School Success. American Educational Research Journal **19**(3) (January 1982) 325–340