Who's playing well with others: Determining collegiality in text

David B. Bracewell, Marc Tomlinson, Sanda Harabagiu, Ying Shi, Jeremy Bensley and Mary Draper
Language Computer Corporation
Richardson, Texas 75080
Email:{david,marc,ying,jeremy,mary}@languagecomputer.com

***Abstract*—In this paper, we present a framework for determining the interpersonal relations exhibited between two individuals. Specifically, we focus on recognizing the presence or absence of collegiality in discussion threads and dialogues. Collegiality results from the existence of harmonious relationships irrespective of the group's power structure. We have identified four psychologically-motived language uses that indicate collegiality. These language uses are identified in text with the use of a set of attributes that are assigned to each language use and can be extracted using grammars and lexicons. Through the attributes, language uses, and dialogue features, a model can be learned that can determine whether two people are collegial, uncollegial, or whether there is not enough information. Using multi-class logistic regression, we obtain an overall micro-averaged F-measure of 83.3%.**

I. Introduction

Humans are social animals, and their lives are made up of numerous interactions between individuals. These repeated interactions can be simple by being either friendly or hostile, or they can be complex and filled with subtle nuances. These interactions define the relationships and social groups within a person's personal and professional life. In this paper, we describe a method of identifying one class of interactions between people, collegiality, in online forums and meeting transcripts. Collegiality is a social construct defined as two or more individuals cooperating in reaching a common goal or ideal. It has an impact on everything from career advancement [1] to team performance [2] and the meaning of life [5].

Collegiality has importance at a personal, interpersonal, and group level. At the personal level, collegiality can be an indicator of the degree of social support a person has. Social support can be physical or emotional and can have effects on job satisfaction [3] and, in some cases, quality of life [4]. It often times comes from friends, family, co-workers, and others that may not normally be considered to be in one's social network. Collegiality can also aid in mapping a person's interpersonal social network information. A person's social network is often made up of mostly friends and family. Collegiality can be a signal of friendliness, and, as such, it may be in an important indicator of who is in a person's social network. In addition, identifying collegiality within a group is important for anticipating group behavior and performance. Team performance is improved when the team is collegial [2]. Collegiality is an important determinant of how likely a team is to reach its goals [6].

In this paper, we look at determining three levels of collegiality: collegial relationships, uncollegial relationships, and relationships where there is Not Enough Information to Decide (NEID). These last cases may contain either relationships where there is an insufficient number of interactions between the individuals or longer relationships where there are no indicators of positive or negative collegiality. The degree of collegiality of each case is determined through the identification of four psychologically and linguistically motivated language uses (LUs) and a number of discourse features and non-LU socio-linguistic features. Identification of the LU is accomplished through extraction of linguistic markers, using a content extraction system. A two-tier system is created in which the first tier models the LUs using the linguistic markers, while the second tier models the social construct using the LUs as well as various discourse features. In this particular research, the model focuses on collegiality, but it can also be extended to apply to other social constructs. Moreover, the approach is language-independent.

The paper will proceed as follows. In section II, related work is examined. Next, in section III, a thorough look at how to model collegiality is offered and the motivations for each of the language uses provided. In the following section, IV, the two-tier model and the learning process are explained. In section V, details of the collected data and experimental results are presented. Finally, in section VI, concluding remarks are given and considerations for future work proposed.

II. Related Work

Work in the area of social relationship extraction can be divided into several areas. The field of socio-linguistics boasts well-established studies of interpersonal relationships. For example, Eggins and Slade present a thorough linguistic analysis on causal conversations that covers topics such as humor, attitude, friendliness, and gossip [7]. This is accomplished through a comprehensive analysis of the dialogue at multiple levels. In contrast, however, research using Natural Language Processing to automatically identify social relationships in text is still in its infancy.

Strzalkowski et al. (2010), using an approach analogous to ours, break down social phenomena into mid-level social language uses. They focus on the use of discourse features (e.g. topic control) to identify language uses [8] that might be indicative of some social constructs. In contrast, we focus on using multiple types of learned LU's, discourse features, and non-LU sociolinguistic features.

Alternatively, network based-approaches have also been gaining popularity in the areas of social roles, events, and groups. Roles, events, and groups can influence and constrain the interpersonal relationships between individuals. Agarwal and Rambow (2010) look at extracting two types of social events in text: interaction and observation. These events are then used in extracting social network models from text. They use Support Vector Machines with tree kernels based on phrase structure and dependency trees [9].

These social roles may influence interpersonal relationships; thus, Brendel and Krawczyk use the relations of people in their social network to determine their roles [10]. Similarly Fisher et al. (2006) examined social networks from Usenet newsgroups. They found that the type of newsgroup determines the type of relationships and roles that are found [11].

Sentiment analysis, which includes polarity and affect classification, is often used in classifying relationships or groups as well. Thomas et al. use Support Vector Machines to identify support and opposition, a sub-task of sentiment analysis, for legislation from transcripts of Congressional debates [12]. Thelwall et al. found that emotion was important in expressing friendship on social networks [13].

Non-verbal features have been used for identifying relationships and social roles in face-to-face interactions. Zancanaro et al. (2006) examine the identification of social roles in face-to-face interactions. They use visual and auditory information from scenes to train a model to detect and classify the social roles present [14]. Gatica-Pereze offers a review of non-verbal analysis of social interactions, bringing together literature from both the computational and social sciences [15].

III. Modeling Collegiality

Collegiality is defined as two or more individuals who cooperate in reaching a common goal or ideal. We define negative collegiality, or uncollegial behavior, as actively sabatoging the work of an individual towards the common-goals of the group. Collegiality between individuals is sometimes directly referenced in text, but more often, we must look for language uses that indicate the status of the relationships between individuals. From literature in the field of psychology, we have identified several types of language uses that are possible indicators of collegiality. The amount of evidence that these attributes provide for establishing collegiality is learned through a multi-tiered regression model.

A. Language Uses

1) Establish Mutually Supportive Behavior: By definition, mutually supportive behavior, or cooperation towards a common goal, is an example of collegiality. This type of behavior lies at the center of group dynamics. Cooperation is correlated with both overall group performance and managerial ratings of group effectiveness [6].

Evidence for cooperation manifests itself in many different ways. Classically, there is the notion of cooperation on a physical task (e.g. one person helping another lift a heavy weight), or cooperation through social support (e.g. Mary says, "John's decision is excellent"). However, care must be taken when looking for evidence of cooperation cross-culturally, as assignment of these attributes to the LU is not always straightforward. Different cultures can exhibit differences in whether or not helping is actually cooperative. For example, in Chinese groups, high levels of helping can be construed as a negative indicator of cooperation (one person is taking advantage of the other), whereas in the United States being helpful is almost always considered a positive indicator of cooperation [16].

There are also more subtle, unconscious examples of cooperation between individuals, which can demonstrate a certain degree of collegiality between the individuals. One example is cooperation for the effective use of language and the building of dialogue [17]. Dialogue is a complicated interaction that requires commitment from both parties. In order to maintain a stable conversation, participants must be willing to expend cognitive effort to listen, understand, and form a relevant response that advances the dialogue. The degree to which participants are able to maintain a cohesive dialogue should be reflected in the collegiality of the participants. If one participant is not cooperating, the dialogue will not progress. Attributes following this approach have the advantage of being domain- and culture-independent.

Six linguistic markers that indicate the 'establish mutually supportive behavior' language use have been identified.

- Ideas / Actions to be supported: An idea that a speaker introduces to the group for support.
- Encouragement: Statement of support directed at another speaker's idea/action or the speaker himself/herself.
- Praise: Positive statement lifting up a speaker or his/her work.
- Intention to help: Statement made by a speaker offering help.
- Recapitulation: Restatement of a speaker's idea/action by a different speaker.
- Additional support for position of support: Mention of a previously unmentioned benefit or strength of idea/action.

2) Acknowledged Shared Values: Similar to cooperation, acknowledgment of shared values is a good indicator of collegiality. In order for two people to cooperate in working towards a common goal, they must share that common goal. Additionally, shared values suggest common group membership, which is indicative of a collegial relationship between the individuals.

Groups are formed by people to accomplish a common goal. Evidence of a common value, or goal, between individuals supports the notion that the two individuals are - or will be - in the same group.

Four linguistic markers that indicate the 'acknowledge shared values' language use have been identified.

- Value: What a speaker considers to be an important end goal.

- Value-Impact: What is necessary in order for the end goal - or value - to be achieved.
- Opinion / Goal: Word or group of evaluative words that show how the value-attribute affects the value, i.e. 'important,' 'necessary,' 'destructive,' etc.
- Reaction from other people: Positive response to a speaker's statement of value on behalf of the entire group.

3) Establish Solidarity: Further, language indicative of a desire for group solidarity encapsulates the establishment and maintenance of shared group membership. Group membership can be expressed at either the relational level (e.g. Father, co-worker, etc.) or the collective level (e.g. single mothers) [18]. This desire to establish solidarity can be expressed explicitly (e.g. "We're all in this together"), covertly (e.g. as through the use of inclusive first-person pronouns), or through unconscious actions and linguistic cues, such as the use of in-group jargon, certain syntactic constructions, and mimicry.

Mimicry - both verbal and physical, semantic and syntactic - plays an important role in demonstrating cohesion between individuals. Individuals are more likely to mimic in-group members than members outside of the group [19]. This use of mimicry extends to shared function word usage between individuals, which ebbs and flows with their relationship status [20]. While physical mimicry is easy to fake, evidenced by the fact that many experiments studying mimicry use confederates that are instructed to mimic the individual, it is an open question as to how difficult artificial mimicry of function word usage is. Function words also have the property of being easily detected across domains and cultures.

- Introduction to group: Speaker identifies him/herself during first time in a group.
- Establish bona fides: Speaker establishes good faith with group by stating good intentions or offering help.
- In group jargon: Speaker uses group-specific words or phrases that have special meanings.
- Disclose personal data: Speaker gives personal information about him/herself to the rest of the group. This often occurs shortly after giving a personal self-introduction and serves the purpose of establishing solidarity with the group.
- Disclose beliefs: Speaker shares his/her belief about something in order to establish solidarity with the rest of the group (often occurs shortly after speakers self-introduction to group)
- Ask for a favor: Speaker asks other members of the group to help him/her out.
- Address fallout/conflict: Speaker addresses a past, present, or potential future conflict within the group and states his/her intention to move beyond it (making peace).
- Identify allies: Speaker identifies an ally common to group members; ally may be inside (must be marginalized) or outside the group.
- Identify opponents: Speaker identifies an opponent common to group members; opponent may be inside or outside the group.

4) Offer Gratitude: Lastly, there is psychological validation for the consideration of attitudes expressed by one individual towards another. Even in the absence of any major differences within a group, the expression of an in-group bias and out-group bias [21] between individuals still takes place. Individuals within a group are more likely to possess positive feelings for another individual within the group and to rate him or her more highly than an individual outside of the group.

Five linguistic markers indicating the 'offer gratitude' language use have been identified.

- Thankfulness in reaction to some action: Speaker makes a simple expression of gratitude.
- Preemptive gratitude: Speaker intentionally and without prompting inserts a statement of gratitude into conversation.
- Cheering: Speaker issues celebratory remarks in response to some other statement.
- Open recognition: Speaker deliberately and publicly recognizes some achievement or praiseworthy action of another speaker.
- Supportive of speaker's points: Speaker restates previous speakers points with supportive/thankful comments after each point.
- Acceptance of thankfulness: Speaker indicates his/her reception of another's statement of gratitude.

B. Other Features

In addition to the aforementioned language uses, a number of discourse level, and non-LU socio-linguistic features are also used in determining collegiality. These features cover the number of turns[1], number of interactions between individuals, number of replies, whether or not the individual started the thread/meeting, word count, function word overlap (evidence of mimicry), and presence of offensive and negative sentiment terms.

- No. of speaker turns: Number of times speaker spoke or commented.
- No. of target turns: Number of times target spoke or commented.
- No. of speaker replies: Number of times speaker replied to any other speaker or commenter.
- No. of target replies: Number of times target replied to any other speaker or commenter.
- No. of interactions: Number of times speaker and target interacted with one another.
- No. of speaker words: Number of words speaker used in all comments/dialogue.
- No. of target words: Number of words target used in all comments/dialogue.
- Function word overlap between speaker and target: Jaccard similarity [22] of function word usage between speaker and target.

[1] A turn is an instance of communication by one speaker, such as a comment, utterance, message, etc.

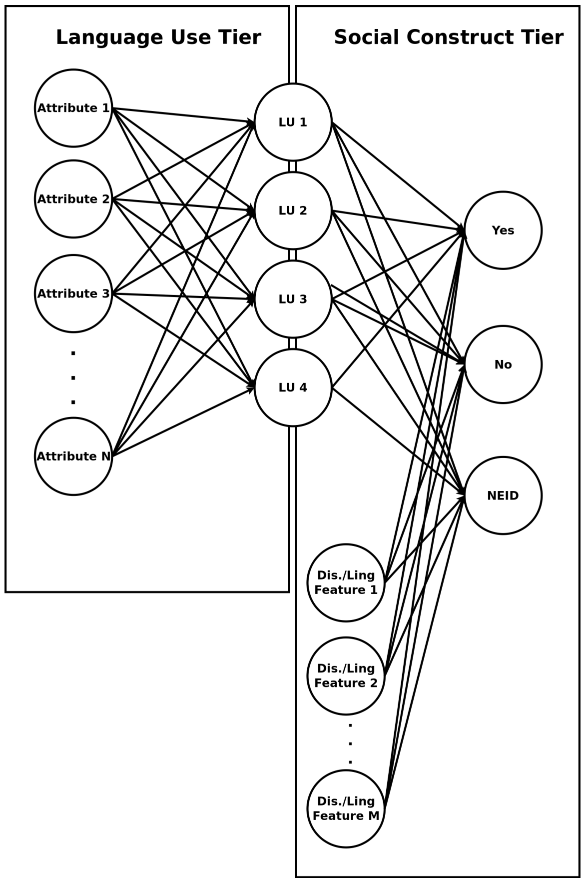

Fig. 1: Multi-Tier Learning Model

- Negative sentiment: Amount of negative sentiment used by speaker.
- LUs per interaction: Average number of LUs per interaction between speaker and target.
- Negative sentiment per interaction: Average amount of negative sentiment per interaction between speaker and target.

IV. Learning

A multi-tier learning framework is adopted as shown in figure 1. The first tier is the language use learning tier where individual langauge uses are identified through the presence and absence of linguistic markers (attributes). The second tier identifies the existence and valence of the social construct, in this case collegiality, based on the the evidence from the indvidual languages as determined by the first tier. Each tier uses logistic regression for classification.

A. Language Uses

Attributes for each turn of the dialogue/thread are first extracted using lexical and grammatical patterns; see section III for descriptions of the attributes associated with each language use. These patterns for the attributes were hand-coded based by members of our team. The extracted attributes are used as input to the language use classifiers, with each attribute either being present or not present for each speaker turn. On average we identified 2.3 attributes per turn.

Each language use is then modeled independently from all of the available attributes using logistic regression. The presence of a language use is determined independently for each turn in the dialogue. Positive outcomes indicate the presence of the corresponding language use with a confidence equal to the probability given by the classifier. The system allows for a correspondence between any attribute and any language use, but we found good consistency in how the attributes aligned with the particular language uses in our training data.

B. Social Construct Learning

Social construct learning is the second tier of learning in the system. It uses multi-class logistic regression to determine values of positive, negative and Not Enough Information to Determine (NEID). The inputs to the classifier are the language uses found in the first tier of learning, discourse metrics, and non-LU specific sociolinguistic features determined from the interactions between two people (target pair).

For social constructs that characterize interpersonal relationships, such as collegiality, only the language uses in turns that represent interactions between the two people for which the social construct is trying to be identified. The input values of the language uses are the sum of the confidences associated with the language uses (output from the first tier). The discourse and non-LU specific sociolinguistic features are used to aid in identification where the language uses alone may not be enough, such as in identifying NEID answers.

C. Example

Figure 2 shows an example from a Wikipedia talk page[2]. The example shows our systems annotation and discovered LUs for a positive instance of collegiality between the authors ”Ex nihil” and ”Nick Dowling”. Turns 2 through 4 and turns 8 and 9 are interactions between the two authors.

V. Experiments and Results

While learning is performed to determine language uses in text from textual attributes, the focus of our experimentation is the final output of the second tier: collegial, uncollegial or Not Enough Information to Determine (NEID). This section will discuss the data that we used as well as the experimental results.

A. Data

We collected data from two sources: Wikipedia talk pages[3] and Nuclear Waste Technical Review Board (NWTRB) meeting transcripts[4]. A combined total of 55 threads/meetings have been collected. Each is annotated with both language uses and collegiality. Collegiality annotation looks at all speaker/author pairs and assigns ’collegial,’ ’uncollegial,’ or ’NEID’.

[2]http://en.wikipedia.org/wiki/Talk:Timor_Leste_Defence_Force#Move_to_East_Timor_Defence_Force

[3]http://en.wikipedia.org/wiki/Wikipedia:Talk_page

[4]http://www.nwtrb.gov/meetings/meetings.html

Turn 1) Propose that this page be moved to East Timor Defence Force as this is the closest translation of Foras de Defesa de Timor Leste. I have worked in Timor Leste as a government advisor, including with FDTL, and have never heard anybody ever refer to the FDTL as Military of East Timor. Ex nihil (talk) 06:57, 5 March 2008 (UTC)

Turn 2) As I understand it, 'East Timor Defence Force' is considered outdated. While it was commonly used when the force was established, almost all english-language publications now use 'F-FDTL'. 'Military of East Timor' is a generic name, and **I agree** [*LU-4 supportive of speaker's points*] that it's rarely used and not a great title. **I'd prefer 'Timor Leste Defence Force'** [*LU-1 ideas/actions*] as this seems to be the direct translation, but this would be inconsistent with the other Wikipedia articles on the country. Should we be bold and move this article to 'Timor Leste Defence Force'? –Nick Dowling (talk) 07:47, 5 March 2008 (UTC)

Turn 3) I so totally agree with you [*LU-1 encouragement*]. 'Timor Leste Defence Force' is it. The only reason I did not propose that was the failure to change the country page from East Timor to Timor Leste, a decision that I feel was extremely discourteous of Wikipedia considering the government's specific request that it be referred to as Timor Leste. If you have worked there you will know that everybody uses 'Timor Leste', even the ADF but the Australian DFAT uses East Timor although the more enlightened Kiwi embassy call it TL. **I suggest we leave it for 48 hours to see if anyone has any strong feelings and then change it to ' Timor Leste Defence Force' with diverts from F-FDTL and FDTL** [*LU-1 ideas/actions*]. Ex nihil (talk) 08:41, 5 March 2008 (UTC)

Turn 4) I agree with that approach [*LU-4 supportive of speaker's points*]. In the interests of consensus editing, I've posted a note at Talk:East Timor (in lieu of a Wikiproject on the country) to seek other editors' views. –Nick Dowling (talk) 09:54, 5 March 2008 (UTC)

...

Turn 8) As no-one has raised any objections, I've just made the move. –Nick Dowling (talk) 07:51, 9 March 2008 (UTC)

Turn 9) Good move, well done [*LU-4 praise*] Ex nihil (talk) 12:07, 9 March 2008 (UTC)

Fig. 2: Example Dialogue from Wikipedia

Each Wikipedia talk page is associated with one Wikipedia article. The talk pages are made up of threaded comments and provide a place for users to discuss the target article and propose changes to its text. A collection of 14 talk pages that covered a wide domain of topics was used. The talk pages were split into 26 comment threads. On average, each thread was found to contain 6.2 comments, 3.8 authors 2.2 attributes, and 2.6 language uses.

The NWTRB is an independent agency providing oversight to the Department of Energy in regards to nuclear waste management and disposal. Transcripts of its meetings are available on its website for download. The meetings consist of presentations and discussions by various board members, as well as a question and answer period from the general public. A collection of 29 meetings ranging from May 2000 to October 2010 was collected. On average, each meeting was shown to have 179.6 turns, 15 speakers, 2.3 attributes, and 2.3 language uses.

Between both sets of documents, the annotators found a total of 50 pairs of speakers/authors who are collegial and 20 who are uncollegial. This left a remaining of 4,669 speaker/author pairs for which there was not enough information to determine collegiality.

TABLE I: Results for Collegiality Classification

	Precision	Recall	F-Measure
Collegial	82.7%	86.0%	84.3%
Uncollegial	62.5%	50.0%	55.6%
NEID	90.4%	94.0%	92.2%
Micro-Avg.	83.3%	83.3%	83.3%
Macro-Avg.	78.5%	76.7%	77.4%

Due to the unbalanced nature of the data, resampling was performed to balance the class distribution. First, 50 random NEID examples were chosen from the data to create a 120 item data set. This data set was then used for evaluation. The system was evaluated using 'leave one out' cross-validation. During each fold SMOTE [23] resampling was used to balance the class proportions.

B. Results

For testing collegiality, 'leave one out' cross-validation was performed. The data was resampled during each turn to make the training set uniform. Table I shows the results.

The resultant micro-average F-Measure is 83.3% and macro-average F-Measure is 77.4%. The positive (collegial) class, which the language uses are designed to identify, had an F-Measure of 84.3% with a slightly higher recall than precision. The discourse features aid in accurately determining the NEID class by indicating whether or not two speakers interacted at all or had enough substantial interactions. The negative (uncollegial) class had the fewest examples and performed the worst. The only features that were specifically geared toward identifying it are Negative sentiment and Negative sentiment per interaction.

VI. Conclusion

In this paper, we presented a two-tier framework for identifying collegiality in text using language uses, discourse features, and sociolinguistic cues. The first tier was made up of a set of a four binary classifiers that identified the presence or absence of particular language uses. The nature of these language uses was motivated by work in psychology and linguistics. They were identified by the presence of linguistic, semantic, and pragmatic markers in the text. The second tier of the system used the output from the first tier, combined with additional discourse and sociolinguistic features, to train a three class model that determined collegiality, uncollegial or Not Enough Information to Decide. The second tier was able to achieve a micro-averaged F-measure of 83.3%.

Overall, the system did a very good job at identifying the presence of collegiality between individuals; however, the system was not able to capture uncollegial behavior very well. We believe that this is due to the data sets that were used. One property of professional meetings and even (perhaps surprisingly) Wiki-talk data is that most individuals are actually fairly nice, generally well-behaved, and respectful toward each other. We feel that additional data sets such as forums that have a broader use should yield data that support easier identification of uncollegial behavior.

In the future, we plan to extend this general methodology to look at other social constructs, such as acquiring power. The acquisition of power is particularly important for individuals in the Wiki-talk discussions because they aim to sway the opinions of other editors.

We also plan to add a feedback mechanism to the language use classifiers that is based on the performance of the social construct classifier. One possible implementation of this would be to simply increase the weight of LU examples belonging to social constructs that were misclassified. Another option would be to change the learning framework to use a learning algorithm that has explicit feedback built in, such as a neural network.

ACKNOWLEDGMENT

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] B. Balsmeyer, K. Haubrich, and C. Quinn, "Defining collegiality within the academic setting." *The Journal of nursing education*, vol. 35, no. 6, pp. 264–267, 1996.

[2] J. W. Little, "Norms of Collegiality and Experimentation: Workplace Conditions of School Success," *American Educational Research Journal*, vol. 19, no. 3, pp. 325–340, Jan. 1982. [Online]. Available: http://aer.sagepub.com/cgi/content/abstract/19/3/325

[3] B. S. McCann, J. Russo, and G. A. Benjamin, "Hostility, social support, and perceptions of work." *Journal of occupational health psychology*, vol. 2, no. 2, pp. 175–85, Apr. 1997. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/9552289

[4] S. L. Shapiro, A. M. Lopez, G. E. Schwartz, R. Bootzin, A. J. Figueredo, C. J. Braden, and S. F. Kurker, "Quality of life and breast cancer: relationship to psychosocial variables." *Journal of Clinical Psychology*, vol. 57, no. 4, pp. 501–519, 2001. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/11255204

[5] T. F. Stillman, R. F. Baumeister, N. M. Lambert, a. W. Crescioni, C. N. Dewall, and F. D. Fincham, "Alone and Without Purpose: Life Loses Meaning Following Social Exclusion." *Journal of experimental social psychology*, vol. 45, no. 4, pp. 686–694, Jul. 2009. [Online]. Available: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2717555\&tool=pmcentrez\&rendertype=abstract

[6] M. Campion, E. Papper, and G. Medsker, "Relations between work team characteristics and effectiveness: A replication and extension," *Personnel psychology*, vol. 49, no. 2, pp. 429–452, 1996. [Online]. Available: http://onlinelibrary.wiley.com/doi/10.1111/j.1744-6570.1996.tb01806.x/abstract

[7] S. Eggins and D. Slade, *Analysing casual conversation*. Cassell, 1997.

[8] T. Strzalkowski, G. A. Broadwell, J. Stromer-galley, S. Shaikh, S. Taylor, and N. Webb, "Modeling Socio-Cultural Phenomena in Discourse," in *23rd International Conference on Computational Linguistics*, no. August, Beijing, 2010, pp. 1038–1046. [Online]. Available: http://www.aclweb.org/anthology/C/C10/C10-1117.pdf

[9] A. Agarwal and O. Rambow, "Automatic detection and classification of social events," in *EMNLP '10 Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, Stroudsburg, Oct. 2010, pp. 1024–1034. [Online]. Available: http://portal.acm.org/citation.cfm?id=1870658.1870758

[10] R. Brendel and H. Krawczyk, "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, pp. 163–170. [Online]. Available: http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=4573053

[11] D. Fisher, M. Smith, and H. T. Welser, "You Are Who You Talk To: Detecting Roles in Usenet Newsgroups," *Proceedings of the 39th Annual Hawaii International Conference on System Sciences HICSS06*, vol. 00, no. C, pp. 59b–59b, 2006. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1579410

[12] M. Thomas, B. Pang, and L. Lee, "Get out the vote: determining support or opposition from congressional floor-debate transcripts," in *EMNLP '06 Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, Jul. 2006, pp. 327–335. [Online]. Available: http://portal.acm.org/citation.cfm?id=1610075.1610122

[13] M. Thelwall, K. Buckley, G. Paltoglou, and D. Cai, "Sentiment Strength Detection in Short Informal Text," *Journal of the American Society for Information Science*, vol. 61, no. 12, pp. 2544–2558, 2010. [Online]. Available: http://www.scit.wlv.ac.uk/\~{}cm1993/papers/SentiStrengthPreprint.doc

[14] M. Zancanaro, B. Lepri, and F. Pianesi, *Automatic detection of group functional roles in face to face interactions*. New York, New York, USA: ACM Press, Nov. 2006. [Online]. Available: http://portal.acm.org/citation.cfm?id=1180995.1181003

[15] D. Gatica-Perez, "Automatic nonverbal analysis of social interaction in small groups: A review," *Image and Vision Computing*, vol. 27, no. 12, pp. 1775–1787, Nov. 2009. [Online]. Available: http://dx.doi.org/10.1016/j.imavis.2009.01.004

[16] J. Keller and J. Loewenstein, "The Cultural Category of Cooperation: A Cultural Consensus Model Analysis for China and the United States," *Organization Science*, vol. 22, no. 2, pp. 299–319, Jun. 2010. [Online]. Available: http://orgsci.journal.informs.org/cgi/doi/10.1287/orsc.1100.0530

[17] S. Garrod and M. J. Pickering, "Why is conversation so easy?" *Trends in Cognitive Sciences*, vol. 8, no. 1, pp. 8–11, Jan. 2004. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S136466130300295X

[18] M. B. Brewer and W. Gardner, "Who is this "We"? Levels of collective identity and self representations." *Journal of Personality and Social Psychology*, vol. 71, no. 1, pp. 83–93, 1996. [Online]. Available: http://doi.apa.org/getdoi.cfm?doi=10.1037/0022-3514.71.1.83

[19] Y. Yabar, L. Johnston, L. Miles, and V. Peace, "Implicit Behavioral Mimicry: Investigating the Impact of Group Membership," *Journal of Nonverbal Behavior*, vol. 30, no. 3, pp. 97–113, 2006. [Online]. Available: http://www.springerlink.com/index/10.1007/s10919-006-0010-6

[20] M. E. Ireland and J. W. Pennebaker, "Language style matching in writing: synchrony in essays, correspondence, and poetry." *Journal of personality and social psychology*, vol. 99, no. 3, pp. 549–71, Sep. 2010. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/20804263

[21] M. B. Brewer, "In-group bias in the minimal intergroup situation: A cognitive-motivational analysis." *Psychological Bulletin*, vol. 86, no. 2, pp. 307–324, 1979. [Online]. Available: http://content.apa.org/journals/bul/86/2/307

[22] P. Jaccard, "Étude comparative de la distribution florale dans une portion des Alpes et des Jura," *Bulletin del la Société Vaudoise des Sciences Naturelles*, vol. 37, pp. 547 – 579, 1901.

[23] K. W. B. Nitesh V. Chawla, "SMOTE: Synthetic Minority Over-sampling Technique," *Journal of Artificial Intelligence Research*, vol. 16, pp. 321–357, 2002.