

The Language of Power and its Cultural Influence

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 goals of an individual can be recognized through analysis of the social actions indicated by their use of language. Specifically, we focus on recognizing when someone is pursuing power within a web forum. Individuals pursue power in order to increase their control over the actions and goals of the group. We cast the problem as social conversational entailment where we determine if a dialogue entails a hypothesis which states a dialogue participant is in pursuit of power. In the social conversational entailment framework the hypothesis is decomposed into a series of social commitments which define series of actions and responses that are indicative of the hypothesis. The social commitments are modeled as social acts which are pragmatic speech acts. We identify nine culturally neutral psychologically-motivated social acts that can be detected in language and are indicative of whether an individual is pursuing power. Our best results using social conversational entailment achieve an overall F-measure of 79.7% for predicting pursuit of power for English speakers and 78.3% for Chinese speakers.

Keywords: dialogue, power, social actions, entailment, online communication, culture, norms.

1 Introduction

Social media has empowered the masses by allowing individuals to participate in a variety of group projects which impact the future of society. Sites like Wikipedia allow anyone to edit content that is used the world over for resolving debates and informing science. Because of the influence of these sites, many contributors pursue a high-power role giving them control over the site's content and the goals and actions of other contributors. Traditionally research has focused on inferring whether individuals are already in power through such means as social network analysis or more recently through the language those in power employ (Bramsen et al., 2011; Danescu-Niculescu-Mizil et al., 2012). However, the dialogue taking place on these sites provide a richer source of observations on the interaction patterns of individuals and in particular how individuals pursue power. By analyzing and characterizing these observations, models of these pursuits of power can be built which can provide important information about the dynamics of the group and its evolving leadership structure. To accurately infer if an individual is pursuing power, we must address three main questions:

- (1) *What characterizes the language of individuals pursuing power?*
- (2) *Can these characterizations be captured automatically?*
- (3) *What impact does culture have on the characterization of pursuit of power?*

In order to answer these questions, we must first define what it means for an individual to pursue power. Power is a nebulous topic whose meaning differs depending on the domain and the context. For the purposes of this paper, we have defined the pursuit of power for online discussions, such as those on Wikipedia, as repeated attempts by an individual to increase their status and to control the actions and goals of others. As an example let us examine the dialogue in Figure 1 which illustrates a pursuit of power by participant B.

A) B ... contrary to your edit summary, your deletion of over 4,000 bytes of material is not explained on the talk page. Please cease and desist. It is vandalism to delete sourced information relevant to the topic at hand. As an editor states directly above, it is more beneficial to add to the article than to delete material that is reliably sourced. Please do so.

B) The explanation is right above - read it, and stop claiming this is vandalism when it is not.

A) There is no explanation from you above, only (with regret) some hyperbole. As the editor writes above: "It goes without saying that examples of use should be included. It can be expanded later. Add more material rather than deleting." Please cease deleting this material without consensus.

B) Please read more carefully. The explanation is there. Your repeated claims of vandalism are uncivil - stop them.

C) I tend to agree with B this article is not from a neutral point of view.

Figure 1: Example conversation in which participant B is pursuing power.

As the dialogue in Figure 1 illustrates the understanding of social phenomena is less about the content of the information exchanged and more about the social actions and intentions of the participants. Building on the work done in textual entailment (Bar-Haim et al., 2006; Dagan et al., 2005; Giampiccolo et al., 2007; Hickl, 2008) and conversation entailment (Zhang and Chai, 2010), we cast the problem of implying the social phenomena, e.g. pursuit of power, exhibited by participants in a dialogue as *social conversational entailment*. Textual entailment has focused mainly on the information, often factual, distilled in monologue. Conversation entailment extended this to the information exchanged in conversation. In contrast, social conversational entailment focuses on the social phenomena exhibited by the participants in the conversation through examination of their social intentions and goals which are captured through social acts. The intentions and goals informs the *why* and constrains the *how* information is exchanged. For example, in Figure 1 the information

exchanged by B was to support his desire to gain power and control the content of the Wikipedia article. Additionally, B 's actions influenced A to communicate in such a way as to try to stave off B 's potential gain in power.

In this paper, we explore building social conversational entailment models to entail hypotheses about whether or not an individual is pursuing power. We focus on groups communicating in English and Chinese on Wikipedia talk forums. We then turn our focus to how cultural differences exhibit themselves in the characterization of pursuit of power.

2 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 (1997) present a thorough linguistic analysis on causal conversations that covers topics such as humor, attitude, friendliness, and gossip. Other studies have examined how individuals vie for power in meetings and the work place (Keller, 2009; Owens and Sutton, 2001). These studies have shown that status differences can have a large effect on how a particular individual will seek power. Further, work on the effects of power on cognition has shown that individuals with power use language differently than lower status individuals (Smith and Trope, 2006) has provided insights on how pursuits of power may be characterized.

Recently work in natural language processing has been conducted to identify the relative status of individuals through automated analysis of their language. Bransen et al. (2011) looked for the presence of *upspeak* (speech directed towards individuals of higher status) and *downspeak* (speech directed towards individuals of lower status) within the Enron email corpus using an n-gram based approach combined with human-engineered features. They achieved an accuracy of 78.1% for detecting the relative status difference between individuals. Additionally, Danescu-Niculescu-Mizil et al. (2012) examined the use of coordination, often referred to as mimicry, for inferring power relationships. In contrast to identifying a static social relationship between individuals, we look at detecting an individual's intentions to manipulate an existing social relationship.

There is a long history of work in discourse understanding that focuses on understanding the pragmatics of the discourse. More recent work has focused on inferring information, such as conversational intent, about the discourse participants. Zhang and Chai (2010) introduced conversation entailment, which is designed to answer a variety of hypotheses about dialogue participants. The hypotheses can be about factual information, beliefs and opinions, desires, or communicative intentions. Additionally, work in textual entailment has used discourse commitments, which are general beliefs held by the author of the text and hypothesis (Hickl, 2008). In contrast, we focus on social conversational entailment, which uses social commitments, for inferring the social roles, relationships, and intentions of dialogue participants through analysis of their social acts as signaled through the dialogue.

3 Social Conversational Entailment for Pursuit of Power

Power is exhibited in many forms, through physical intimidation, wealth (money, physical resources, or knowledge), or position within a hierarchy. There are variety of methods to pursue power. Moreover, because of the shear variety of methods to pursue power, it is difficult to develop a robust cross-domain text-based recognition approach to identify those who are in pursuit. Instead, we focus on detecting differences in the way people use language

when they are attempting to pursue power. We look to mimic human understanding of power and follow the non-conscious cues provided within a dialogue. We model social conversational entailment for pursuits of power after work in speech act recognition (Stolcke et al., 1998) and language modeling (Niederhoffer and Pennebaker, 2002) using *social conversational entailment*. The task of social conversational entailment is defined as follows:

Given a dialogue (D) and a hypothesis (H) about one or more participants, the goal is to determine if D entails H.

Hypotheses in social conversational entailment describe the *role* (e.g. leader) or *action* (e.g. pursuing power) of a participant whom we label the central individual or the *relationship* (e.g. collegial) between two or more participants whom we label the central group. In this paper, we focus solely on the action of pursuing power for a single central individual. A pursuit of power hypothesis is in the form of: **Person A is pursuing power.**

A pursuit of power hypothesis is decomposable into a number of social commitments. These social commitments represent patterns of action that individuals pursuing power are likely to perform as well as the responses elicited by those actions from others towards those in pursuit. We capture the actions performed by participants as social acts. Social acts are pragmatic speech acts that signal a dialogue participant’s social intentions. The social acts used to identify pursuits of power are discussed in section 4.

The model for social conversational entailment is based on the social commitments and social actions. More formally, given a dialogue D which is represented as a series of social acts s_1, \dots, s_m , performed by the central individual and others directed toward the central individual, and a hypothesis H which is represented as a number of social commitments c_1, \dots, c_n , the prediction of whether D entails H is approximated as:

$$\begin{aligned} P(D \models H \mid D, H) &= P(D \models c_1, \dots, c_n \mid D, c_1, \dots, c_n) \\ &= \prod_{i=1}^n P(D \models c_i \mid D = s_1, \dots, s_m, c_i) \\ &= \prod_{i=1}^n P(s_1, \dots, s_m \models c_i \mid s_1, \dots, s_m, c_i) \end{aligned}$$

One way in which we can model the social commitments is a Markov process over the social acts. Social commitments then become chains of social actions which represent prominent patterns associated with individuals pursuing power. Assuming the Markov process, we can approximate the probability of D entailing H as:

$$P(D \models H \mid D, H) \propto \prod_{i=1}^m P(s_i \mid s_{i+m-1}, \dots, s_{i-1})$$

We can further simplify the model by making a first order Markov assumption, which results in:

$$P(D \models H \mid D, H) \propto \prod_{i=1}^m P(s_i \mid s_{i-1})$$

The entailment model is then built from a corpora of positive entailments, i.e. where D entails H , using Kneser-Ney smoothing (Chen and Goodman, 1996). Dialogues with probabilities over some threshold τ given the entailment model have a sufficient alignment

with the social commitments to entail the hypothesis. One potential problem is that the model can be overwhelmed by repeated exhibition of social and cultural norms which participants follow through the normal course of a conversation. The conflation of these norms with true social commitments hinder the accuracy of the inference as the norms are not a sign of pursuit of power.

3.1 Social and Cultural Norms

In order to accurately infer social phenomena it is critical to take into account social and cultural norms. It is often through the violations of these norms that social phenomena, such as pursuing power, are witnessed. The accurate depiction of social and cultural norms is an entire field of research upon its own. Instead of completely addressing this complex topic, we look to only roughly determine the norms as portrayed by participants in a corpus.

By building a model around the actions of participants who do not entail the pursuit of power hypothesis, we can capture aspects of the social and cultural norms. We call this the *background model*. The background model is built in the same manner as entailment using the following equation:

$$P(D \not\models H | D, H) \propto \prod_{i=1}^m P(s_i | s_{i-1})$$

The data used for building the model are negative entailment examples, i.e. dialogue D in which the central individual is not pursuing power. As the diversity of genre and amount of conversations used for training the background model increases it will more accurately portray the social and cultural norms. By combining the entailment and background models, we can more accurately model the characteristics of pursuit of power and better infer if a participant is pursuing power.

3.2 Inference


We combine the entailment and background model in order to determine if a dialogue D entails a hypothesis H . We predict D entails H when:

$$\beta_0 + \beta_1 \cdot P(D \models H | D, H) + \beta_2 \cdot P(D \not\models H | D, H) + \beta_3 \cdot \frac{P(D \models H | D, H)}{P(D \not\models H | D, H)} > 0.5$$

where β_0, \dots, β_3 are weights controlling the effect that the entailment and background model have in predicting if an individual is in pursuit of power. The β_0 weight is the bias and acts as a prior on the likelihood of a participant to pursue power in the training data. The weights are learned using a linear regression model over the training data. Examples of entailment are assigned the value of 1 and examples of non-entailment are assigned the value of 0 making the final equation result in a probability.

4 The Social Actions of those who Pursue Power

Because individuals rarely explicitly state their intent to pursue power in text, we must look for reflections of their social intentions through their language. We use social acts which are pragmatic speech acts to capture the dialogue participants' social intentions. Social acts are specifically designed to take into account participants' social cognition which constrains their dialogue facilitating the inference of their social goals from their communication.



(a) Ladder model of the path to power. Social acts on the left are directed towards the individual and those on the right are made by the individual.

Agreement	Statements made to indicate a sharing of view about something another member has said or done.
Challenge Credibility	Attempts to discredit or raise doubt about another group member's qualifications or abilities.
Disagreement	Statements made to indicate differences in view about something another member has said or done.
Establish Credibility	Statements made demonstrate knowledge or personal experience in order to look better in the eyes of the group.
Group Affordance	Use of honorifics and deference to show respect and esteem for another group member.
Leadership Avoidance	Attempts to avoid being in a position of control over the group.
Leadership Traits	Common linguistic signs that a person is in power, such as extroversion and locus of control.
Managerial Influence	Statements made to control the discussion with the goal of increasing sway over the group.
Solidarity	Statements made to strengthen the group's sense of community and unity.

(b) The set of nine social acts that capture social moves by individuals pursuing power.

We base our list of nine social acts on the reciprocal influence model of power developed by Keltner et al. (2008), shown in Figure 2a. The employment of social acts by the central individual and the group facilitate a change in the central individual's level of power within the group. For example, the use of Leadership Traits by the central individual moves her up the ladder, i.e. increasing her level of power, whereas if a member of the group employs Challenge Credibility it lessens the central individual's level of power. The complete set of nine social acts with their definition is shown in Figure 2b. A more in-depth discussion on social acts and these in particular can be found in Bracewell et al. Bracewell et al. (2012).

5 Data Collection

We constructed a corpus consisting of English and Chinese Wikipedia talk pages. Each Wikipedia talk page is a threaded discussion and is associated with a Wikipedia article. The talk pages provide a forum for users to discuss and debate the content of the target article as well as propose, vote, an denounce changes to the content. We collected a total of 149 English and 401 Chinese Wikipedia talk discussions whose associated articles covered a wide domain of topics. Within these discussions there were a total of 778 and 3,476 participants respectively for English and Chinese. Each discussion was annotated by three to five annotators, which included annotation of every individual in the discussion as either pursuing or not pursuing power.

We employed both in-house and Mechanical Turk annotators. Annotator training consisted of the definition for pursuit of power and example questions to test understanding of the definition. The Mechanical Turk annotators were further tested to judge their language ability. An annotation was said to have agreement when all or all but one annotator chose the same answer, i.e. 2 out of 3, 3 out of 4, or 4 out of 5 chose yes for pursuit of power. For English, we had agreement rates of 76.0% for our in-house annotators, 67.5% for our Mechanical Turk annotators, and 70.0% combined. For Chinese, we had agreement rates of 85.6% for our in-house annotators, 80.0% for our Mechanical Turk annotators, and 82.8% combined.

6 Experimental Results

For experimentation, we used a standard 80/20 split over the data discussed in section 5, where 80% of the participants were used for training and 20% of the participants were used for testing. We focused our experiments to determine the validity of the social conversational entailment model and of using social acts over a purely lexical approach.

As an alternative to the social conversational entailment model, we examined the use of a Support Vector Machine (SVM) classifier using a linear kernel. SVMs have shown promise for such related tasks as the recognition of dialogue acts (Hu et al., 2009) and the identification of social status (Bramsen et al., 2011). We compared the effectiveness of the social acts for inferring pursuits of power to a purely lexical approach. For the SVM model we extracted n-grams from the utterances of the central individual¹. We pruned the list of n-grams using information gain. We tested with different size n-grams, but report here only the best results which were obtained using a combination of unigrams and bigrams. For conversational entailment, a text was generated based on the utterances of the central individual and others in the group responding to the central individual. The origin, i.e. central individual or other, was denoted using a special symbol prepended to the words. For both models punctuation and symbols were removed and cardinals and proper nouns replaced with generic tags (<CARDINAL> AND <PROPERNOUN>). The results of the experiments are presented in Table 1 for English and Table 2 for Chinese.

English	SVM		SCE	
	N-Gram	Social Acts	N-Gram	Social Acts
Pursuing Power	66.2%	79.6%	53.1%	81.4%
Not Pursuing Power	72.7%	63.6%	64.7%	77.8%
Micro-Avg.	69.8%	73.8%	59.7%	79.7%

Table 1: Resulting F-measure for entailing pursuits of power in English using support vector machines (SVM) and social conversational entailment (SCE) with either word-based n-grams (N-Gram) or social acts as features.

Chinese	SVM		SCE	
	N-Gram	Social Acts	N-Gram	Social Acts
Pursuing Power	42.7%	87.2%	1.1%	75.6%
Not Pursuing Power	60.2%	78.8%	28.0%	80.6%
Micro-Avg.	53.0%	84.0%	16.6%	78.3%

Table 2: Resulting F-measure for entailing pursuits of power in Chinese using support vector machines (SVM) and social conversational entailment (SCE) with either word-based n-grams (N-Gram) or social acts as features.

As can be seen in Tables 1 and 2 using social acts performed better than n-grams for SVM and social conversational entailment. This suggests that social acts capture an intermediate-level concept between words and the social phenomena which provide better evidence for entailing pursuit of power. Chinese saw the biggest boost where the use of social acts brought increases in F-measure of 31% and 61.7% respectively for SVM and social conversational entailment. For English the use of social acts brought increases of 4% for the SVM and 20% for the social conversational entailment model. For both models n-grams worked better for inferring pursuits of power in English than for pursuits of power in Chinese (similar findings are seen in text categorization, see Suzuki et al. (2010)).

¹We tried to incorporate information from the other speakers who were replying to the central individual, as is done with social acts, but this resulted in an inability to identify any positive instances of pursuit of power.

7 Discussion

The pursuit of power is a social construct that embodies significant cultural differences, and thus it is exhibited differently across cultures and languages. In order to judge the cultural impact on pursuits of power in Wikipedia discussions, we examined the differences in social acts. The first social act by individuals in pursuit of power is strikingly different between the Wikitalk discussions in Chinese and those in English. For discussions in Chinese an individual who starts the conversation with group affordance, such as honorifics and respectful sentiments, is most likely pursuing power. Also of interest is that Leadership Avoidance is seen as more likely to entail pursuit of power in Chinese and English. The difference, however, is in how Leadership Avoidance is employed with respects to other social acts. In discussions communicated Chinese, an individual normally exhibits Leadership Avoidance after establishing credibility or through a managerial act. Both of these previous social acts are generally strong indicators that an individual is not pursuing power, however the act of Leadership Avoidance, often manifested through order negation, makes it more likely that the individual is pursuing power.

While we cannot draw any strong conclusions on the exact path individuals follow when pursuing power or make overarching statements about the cultural differences, we can state that there are clearly differences in pursuing power in Wikipedia between groups communicating in English and Groups communicating in Chinese. We leave for future research a deeper study on how to accurately capture these differences through improvements in social act identification and the social conversational entailment model. The cultural differences need not be across language, but also exist within a single language, e.g. mainland China vs. Taiwan. By capturing these cultural differences, we believe we can improve the social conversational entailment model as we can better identify the social cultural norms for each individual in the dialogue.

Conclusion

We have shown that is possible to model pursuits of power by individuals in Wikipedia discussions using social conversational entailment. Social conversational entailment answers hypotheses around social roles, relationships, and intentions of individuals in a dialogue. The entailment is validated by fulfilling social commitments, which are culturally dependent mappings of social acts onto social phenomena, such as pursuit of power. The social acts are pragmatic speech acts that capture the social cognition of dialogue participants and are detected through language usage. We have studied the cultural differences in how pursuit of power is exhibited in English and Chinese Wikipedia discussions. We have found that the entailment models of pursuit of powers differ greatly between the two cultures.

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

- Bar-Haim, R., Dagan, I., Dolan, B., Ferro, L., Giampiccolo, D., Magnini, B., and Szpektor, I. (2006). The second PASCAL recognising textual entailment challenge. In *Proceedings of the Second PASCAL Challenges Workshop on Recognising Textual Entailment*.
- Bracewell, D. B., Tomlinson, M., Brunson, M., Plymale, J., Bracewell, J., and Boerger, D. (2012). Annotation of adversarial and collegial social actions in discourse. In *The 6th Linguistics Annotation Workshop (LAW VI)*.
- Bramsen, P., Escobar-Molana, M., Patel, A., and Alonso, R. (2011). Extracting social power relationships from natural language. *Proceedings of ACL HLT*, pages 773–782.
- Chen, S. F. and Goodman, J. (1996). An empirical study of smoothing techniques for language modeling. In *Proceedings of the 34th annual meeting on Association for Computational Linguistics*, ACL '96, pages 310–318, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Dagan, I., Glickman, O., and Magnini, B. (2005). The pascal recognising textual entailment challenge. In *MLCW*, pages 177–190.
- Danescu-Niculescu-Mizil, C., Lee, L., Pang, B., and Kleinberg, J. (2012). Echoes of power: language effects and power differences in social interaction. In *Proceedings of the 21st international conference on World Wide Web, WWW '12*, pages 699–708, New York, NY, USA. ACM.
- Eggs, S. and Slade, D. (1997). *Analysing casual conversation*. Cassell.
- Giampiccolo, D., Magnini, B., Dagan, I., and Dolan, B. (2007). The third pascal recognizing textual entailment challenge. In *In Proceedings of the ACLPASCAL Workshop on Textual Entailment and*.
- Hickl, A. (2008). Using discourse commitments to recognize textual entailment. *Proceedings of the 22nd International Conference on Computational Linguistics - COLING '08*, (August):337–344.
- Hu, J., Passonneau, R. J., and Rambow, O. (2009). Contrasting the interaction structure of an email and a telephone corpus: a machine learning approach to annotation of dialogue function units. In *Proceedings of the SIGDIAL 2009 Conference: The 10th Annual Meeting of the Special Interest Group on Discourse and Dialogue, SIGDIAL '09*, pages 357–366, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Keller, K. (2009). *Power Conflict: Struggles for Intragroup Control and Dominance*. PhD thesis, University of Maryland.
- Keltner, D., Van Kleef, G. A., Chen, S., and Kraus, M. W. (2008). A reciprocal influence model of social power: Emerging principles and lines of inquiry. *Advances in experimental social psychology*, 40:151–192.
- Niederhoffer, K. G. and Pennebaker, J. W. (2002). Linguistic Style Matching in Social Interaction. *Journal of Language and Social Psychology*, 21(4):337–360.

Owens, D. and Sutton, R. (2001). Status contests in meetings: Negotiating the informal order. *Groups at work: Theory and research*, 14:299–316.

Smith, P. K. and Trope, Y. (2006). You focus on the forest when you're in charge of the trees: power priming and abstract information processing. *Journal of personality and social psychology*, 90(4):578–96.

Stolcke, A., Shriberg, E., Bates, R., Coccaro, N., Jurafsky, D., Martin, R., Meteer, M., Ries, K., Taylor, P., and Ess-Dykema, C. V. (1998). Dialog Act Modeling for Conversational Speech. In *Applying Machine Learning to Discourse Processing*, pages 98–105. AAAI Press.

Suzuki, M., Yamagishi, N., Tsai, Y.-C., Ishida, T., and Goto, M. (2010). English and taiwanese text categorization using n-gram based on vector space model. In *ISITA'10*, pages 106–111.

Zhang, C. and Chai, J. (2010). Towards conversation entailment: an empirical investigation. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, number October, pages 756–766. Association for Computational Linguistics.