Identification of Social Acts in Dialogue

David B. Bracewell,   Marc T. Tomlinson,   and Hui Wang
Language Computer Corporation
Richardson, TX
`{david,marc,hui}@languagecomputer.com`

Abstract

The emergence of dialogue on social medial neccessitates the development of new dialogue processing models. We argue that to address coherence and to infer the implicatures of social dialogue it is vital to understand the social aspirations of the dialogue participants. One key aspect of understanding social dialogue is to understand the intentions and goals of participants. In this paper, we present 11 social acts that capture a broad number of social intentions and goals. We define social acts as pragmatic speech acts designed to give insight into the socio-cognitive processes that individuals unconsciously go through when communicating in dialogue. Identification of the social acts is done using a combination of a generative model in which utterances are generated from gappy patterns, which define a given social act, and a series of binary classifiers. Our experimentation shows that we can capture these social acts with an overall F-measure of 50.4%.

Proceedings of COLING 2012: Technical Papers, pages 375–390,
COLING 2012, Mumbai, December 2012.

375

1 Introduction

Traditional approaches to dialogue processing have been primarily focused on task-based (Grosz, 1978; Traum and Hinkelman, 1992) dialogue. In addition, the recognition of speech acts has proved useful for identifying the structure of the dialogue which takes place in formal meetings (Shriberg et al., 2004) where the dialogue is often a function of the job or position of the participants. Theories of the coherence of discourse and discourse relations (Barzilay and Lapata, 2005; Byron and Stent, 1998; Hobbs, 1979, 1985; Mann and Thompson, 1988; Marcu and Echihabi, 2002) have proved useful for the semantic interpretation of discourse. However, in the world of Twitter, Facebook, and other social media where people voluntarily join in the conversation, dialogue is often focused on the social engagements between participants. These *social dialogues* are often not driven by a common group task, goal or purpose, but by the social aspirations of the participants. As an example let us examine the following excerpt of dialogue from a political debate forum:

A) Seriously, how can anybody still support this president's economic policies? How can HE just continue to do the same things and why can't HE take responsibility for the lousy economy???

B) Not to worry, last time we had a democrat president that was this bad, we had a huge victory!

C) Being neither a Republican or Democrat I wonder if a Republican administration would have done any better? We know McCain's proposed tax and spending policies during the 2008 election would have led to even larger deficits than what Obama was proposing at the time. Not to mention the bailouts and TARP of Bush.

D)What about the showing that when Obama took office, the economy was spiraling down with -9% real rate, losing 700,000+ jobs a month, skyrocketing unemployment, and the stock markets were crashing in the worst recession in 80 years. Not to mention the destruction of the housing market. Oh yeah and we were headed straight for a depression.

In the example shown above, the participants do not form a unified group working toward a common goal or task, but are instead splintered into subgroups which have their own agenda. The dialogue only progresses and stays coherent because the participants wish to further their own goals, e.g. further their bond with others in their subgroup, demonstrate their opposition to other groups, and influence the undecided. These social goals represent an individual's task in a task-oriented dialogue. Individuals construct a plan to accomplish their task and carry out their plan through social actions.

In order to address the coherence and to infer social implicatures from social dialogue it is vital to determine the social intentions and goals of the dialogue participants. The representation of social dialogue plays an important role in facilitating the inference of social intentions and goals. We believe the seminal work on attentions, intentions and the structure of discourse by Grosz and Sidner (1986) is best suited for the inference of social goals. A central component of this approach is the intentional structure, which takes into account the purpose of discourse segments. In social dialogue the purpose of a discourse segment is to further the social goal of the dialogue participants. Thus, by understanding the discourse segments, social goals can be inferred.

A straightforward approach to using Grosz and Sidner's theory for inferring social implicatures is to use the prevailing methods in dialogue processing. Topic modeling, such as Blei et al. (2003), can be used for identification of linguistic structure where topic shifts (Cassell et al., 2001) break the larger dialogue into dialogue segments. Dialogue acts (Allen and

Core, 1997; Stolcke et al., 1998; Bunt et al., 2010), which inform the intentions of dialogue, can be employed to infer the social goals of the participants. The attentional state can be captured using local coherence (Barzilay and Lapata, 2005; Byron and Stent, 1998). The social implicatures of the dialogue can then be inferred through the intentional structure, i.e. the social intentions and goals, and the attentional state, i.e. the focus of the dialogue and its participants.

However, our early experimentation revealed that this straightforward approach of using Grosz and Sidner's framework with prevailing dialogue processing techniques fails to capture the complexities of human social interactions and is incapable of reliably inferring the social implicatures of dialogue. The primary factor for this failure is that traditional dialogue acts are not capable of capturing the nuances of the social intentions and goals of the dialogue participants. Instead of focusing solely on the dialogue, we must also focus on the dialogue participants and how their social aspirations constrain their dialogue. Thus, we must look at the *social intentional structure* of dialogue which models the social intentions and goals of the dialogue participants and how the participants perceive the social intentions of others. The question then becomes: *How can the social intentional structure be captured?*

We propose to use *social acts* for inferring the social intentions and goals of dialogue participants which act to define the social intentional structure of the dialogue. Social acts capture the complex social actions individuals signal through their utterances. While most dialogue acts have some social overtones, they fail to adequately interpret the speaker's social goals. In contrast, the definitions of social acts are specifically designed to take into account participants' social cognition which constrains their dialogue facilitating the inference of social intentions and goals from their communication. We identify a set of 11 social acts, listed in section 3, that capture a variety of social goals and intentions. These social acts come from literature in the fields of psychology and organizational behavior and are motivated by work in dialogue understanding.

As with work in dialogue acts, we identify social acts at the utterance level. We employ a generative model for discovering gappy patterns which capture generalized cue phrases for each of the social acts. A gappy pattern consists of one or more words in between which there can exist gaps, or wildcards, which match any word. Each gap has an associated width which determines how many words the gap can match. The gappy patterns are used as features in a binary classifier. Each social act is associated with a classifier which determines if the social act is manifested or not in an utterance.

2 Related Work

Research understanding the intentionality of dialogue and dialogue has a long history. Some of the earliest work in dialogue processing is on speech acts. Speech acts are actions performed by individuals when making an utterance. Austin (1962) formalized the concept of speech acts by separating them into three classes: (1) *locutionary*, (2) *illocutionary*, and (3) *perlocutionary*. Much of the work in speech acts has been focused on illocutionary acts due to the work of Searle (1969).

Dialogue acts are specialized speech acts which include the internal structure, such as grounding and adjacency pairs, of a dialogue. There are a number of schemes for coding dialogue acts, such as DAMSL (Allen and Core, 1997), VERBMOBIL (Jekat et al., 1995), and DIT++ (Bunt et al., 2010). The DAMSL coding scheme defines dialogue acts that are

forward looking, which are extensions of speech acts, and which are backward looking, which relate the utterance to previous utterances. Frameworks like DIT++ have extended the typical coverage of dialogue acts to encompass a boarder set of acts, such as social obligations. However, when dialogue act schemes incorporate socially motivated acts often they do not fully take into account the multitude of purposes, social intentions, and ultimately the social implicatures of these acts. For example, in the statement "get me a cup of coffee", speech acts would focus on identifying the set of actions that would result from the utterance - presumably the target of the utterance physically going to get a cup of coffe for the speaker. In DIT++ the example utterance would most likely be labeled as "Instruct" which is void of any social implicatures resulting from the instruction. In contrast, social acts reflect the social intention of an utterance focusing on the social implicature of the statement, which in the case of the example utterance is that the speaker is indicating their power over the target.

A vast amount of research has been focused on the coherence of dialogue(Barzilay and Lapata, 2005; Byron and Stent, 1998; Hobbs, 1979; Mann and Thompson, 1988). Mann and Thompson Mann and Thompson (1988) introduce Rhetorical Structure Theory (RST), which was originally developed during the study of automatic text generation. They posit that the coherence of a text is attributed to the rhetorical relations between non-overlapping texts called the nucleus and satellite. The definition of the relations are not morphological or syntactic, but instead are focused on function and semantics. Discourse Representation Theory (DRT) (Kamp, 1984) provides a framework for the semantic understanding of discourse. DRT models the cognitive state of the reader, or hearer, of the discourse using discourse representation structures which convert the discourse into a logical form made up of referents and conditions. In social dialogue we must model and understand the speaker's cognitive state, which informs their social intentions and constrains their actions facilitating the progression and coherence of the dialogue.

Grosz and Sidner (1986) posit a structural approach to dialogue understanding where dialogue is broken into three constituents: *linguistic structure*, *intentional structure*, and *attentional state*. The linguistic structure encompasses how utterances combine together to form larger segments of dialogue. The intentional structure is defined using a single dialogue purpose and multiple dialogue segment purposes. The dialogue purpose is the overarching motivation for the dialogue. For social dialogue, the dialogue purpose should infer the social implicatures of the dialogue. Dialogue segment purposes are sub-components of the larger dialogue purpose and define the intention of a single dialogue segment. In social dialogue, one would expect the dialogue segment purpose to relate to the social intentions of the participants. The final structural component, the attentional state is a property of the dialogue and acts to keep track of the current focus of the dialogue. When dealing with social dialogue the attentional state is influenced by the participants and their social intentions.

The inference of social implicatures and identification of social goals through the use specialized social acts has been the focus of recent research. Bramsen et al. (2011) examined how individuals change their language usage depending on the status of the individual with whom they are communicating. In particular, they examine the use of *upspeak* and *downspeak* for signaling power relationships where *upspeak* is a sign that an individual is communicating with someone of higher status and *downspeak* is a sign that an individual is communicating with someone of lesser power. Danescu-Niculescu-Mizil et al. (2012) exam-

ined the use of coordination, often referred to as mimicry, for inferring power relationships. They showed that individuals are more likely to coordinate with individuals of higher status, i.e. those who have more power, than those with lower or equal status. Bracewell et al. (2011) examined a number of social acts for inferring whether two dialogue participants have a collegial relationship.

Other research has focused on the annotation of social acts. Tomlinson et al. (2012) examined the manifestation of a set of social acts in Arabic for inferring pursuits of power by participants. Bracewell et al. (2012) created an annotated corpus of *collegial* and *adversarial* social actions. Bender et al. (2011) created an annotated corpus of social acts relating to *authority claims* and *alignment moves* for determining authority and influence.

Related work is also found in the methods for identifying dialogue acts. Petukhova and Bunt (2011) examined using Bayes Net and Ripper for classification of high level dialogue acts in the DIT++ schema. They reported F-measures ranging from 62% to 95.1% for the AMI meeting corpus. Webb and Ferguson (2010) introduced an automated method for the extraction of cue phrases for identification of dialogue acts. They obtained an identification accuracy of almost 81% on the switchboard corpus when using cue phrases extracted from a portion of the switchboard corpus and an accuracy of almost 71% when using cue phrases extracted from the ICSI-MRDA corpus. Our approach combines the use of cue phrases and classification. We first extract cue phrases in the form of gappy patterns using a generative model. Then the cue phrases are used in a binary classifier which determines if an utterance is a manifestation of the associated social act.

3   Social Acts

Social interaction is one of the primary reasons for dialogue. Even predominantly task oriented dialogue (e.g. "let's go to the movies") has many possible social implications. The most apparent is the expression of a desire to establish or reaffirm a bond between the individuals. In order to reliably infer these social implications, the social intentions and goals for the dialogue participants must be taken into account.

We label the dialogue segment purpose, or the social intentions of an utterance, as social act. Social acts are pragmatic speech acts that signal a dialogue participant's social intentions. There are a number of social goals which a participant may have, including (1) maintaining an existing role, such as being an authority (Mayfield and Rose, 2011), in power (Bramsen et al., 2011), or collegial (Kim and Galstyan, 2010; Bracewell et al., 2011); (2) gaining a new role, such as by pursuing power (Tomlinson et al., 2012); or (3) maintaining or altering the role or status of another individual in the dialogue. Social acts can be signaled by a variety of cue phrases as well as through a dialogue participant's observation or violation of social norms, or expectations of socially appropriate responses.

The set of acts presented in this section have been derived from work in psychology on power, status, and leadership (Anderson et al., 2001; French and Raven, 1959; Keltner et al., 2008; Owens and Sutton, 2001; Smith and Galinsky, 2010), as well as on conflict and cooperation (Brewer and Gardner, 1996; Deutsch, 2011; Jehn and Mannix, 2001). This set of social acts was designed to have broad coverage, but is not to be taken as an exclusive set. Figure 2 gives an example of a dialogue marked with the social acts.

Agreement can act as an affordance to an individual or as a means to establish solidarity between individuals. Likewise disagreement can act as a way of undermining or challenging

Agreement	Statements that a group member makes to indicate that he/she shares the same 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 a group member makes to indicate that he/she does not share the same view about something another member has said or done.
Disrespect	Inappropriate statements that a group member makes to insult another member of the group.
Establish Credibility	Statements that a speaker makes to demonstrate his/her knowledge or personal experience in order to make him/herself look better in the eyes of the group.
Managerial Influence	Statements that a speaker makes to control the discussion with the goal of increasing sway over the group.
Mediation	Attempts made by a group member to resolve a conflict occurring between other group members.
Relationship Conflict	Personal, heated disagreement between individuals.
Solidarity	Statements that a group member makes to strengthen the group's sense of community and unity.
Supportive Behavior	Statements of personal support that one group member makes toward another.
Task Conflict	Disagreement over the manner in which a task is performed or over the outcome of the task.

Figure 1: The set of eleven social acts.

credibility. However, because of the special status of agreement and disagreement we consider them as two separate social acts.

Agreement can be manifested through simple phrases, such as "I agree", through negations of disagreement, such as "I am not disagreeing with you", and through more complex phrases, such as "What Adam says is in principle correct." Similarly, disagreement is manifested through simple "I disagree" phrases as well as negations of agreement, such as "I definitely do not agree with what you said."

Challenging credibility can be used by an individual to lower the status of other group members (Owens and Sutton, 2001). These challenges can be in demands to prove credibility, such as "prove your lies" and aggressive accusing questions, such as "what does that have to do with what we are talking about?". Challenging credibility can also occur through gossip, such as "X doesn't know what he is talking about". This tactic can be used by group members to moderate the power of a leader who has overstepped their boundaries (Keltner et al., 2008).

Disrespected individuals often feel they have been unjustly treated due to the disrespectful action, causing a social imbalance between them and the perpetrator (Miller, 2001). This social imbalance causes a power differential between the two individuals, thus giving the perpetrator power over the individual. Examples of disrespect include "You are a gigantic hypocrite you know that?" and "Do you speak English well?"

Establishing credibility reflects an attempt by an individual to demonstrate their credibility and fitness for leadership (Keltner et al., 2008). Evidence for establishment of credibility manifests itself in many different ways. The most common in our data set is an explicit mention of the individual's credentials, such as "I am a physicist". Alternatively a person can demonstrate their credibility by providing the group with cited information, such as

Figure 2: Social acts tagged for an excerpt of a discussion taken from a Wikipedia Talk page.

"Article 10.5 paragraph 3 says...". Finally an individual can justify their opinion through the use of logic or citation of personally relevant anecdotes.

Managerial influence is used by individuals to signal that they are a leader. Examples of managerial influence include "Can we focus the discussion" and "Are we still trying to find out where the scholarly consensus is on the matter of Lukan authorship?" Figure 2 has a number of examples of managerial influence, such as A proposing to move the page and B giving factual reasons why "Military of East Timor" is an incorrect name for the page.

A person in power often acts as a mediator for disputes between other group members. Mediation itself is an attempt to resolve a conflict occurring between other group members. Individuals performing mediation may already be in position of power. Examples of mediation include "You really need to back off and take a deep breath" and "Let's just all keep calm, yes?"

Relationship conflict is a personal, heated disagreement between individuals (Jehn and Mannix, 2001). Individuals exhibiting relationship conflict are being adversarial. Examples of relationship conflict include "your arrogant blathering" and "I consider it offensive for you to assert that i insist on turning every interaction into a personality conflict.'

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) (Brewer and Gardner, 1996). Language indicative of a desire for group solidarity demonstrates that an individual identifies with the group, an important characteristic of leaders (Keltner et al., 2008) and cooperators (Deutsch, 2011). This 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.

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 (Campion et al., 1996). 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").

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 (Garrod and Pickering, 2004). 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.

Task conflict often arises during power struggles in a group where one individual is attempting to overtake the position of another (Jehn and Mannix, 2001). It is defined as disagreement over the manner in which a task is performed or over the outcome of the task. Task conflict can be manifested by actions performed to undo or challenge other's work toward a task, such as "I reverted your all your changes." Additionally, it may manifest itself as taking sides or stating positions around a conflict, such as "So yes, I will not be editing but I will be monitoring to see if some other naive soul wishes to and try to support them (and revert the vandalism that happens from time to time)."

4   A Generative Model for Identifying Social Acts

A system for the recognition of intentions and goals is the foundation for the understanding of social implicatures in social dialogue. Here we present a generative model for identifying social acts at the utterance level. The model generates utterances exhibiting a social act as a series of gappy patterns. A gappy pattern consists of one or more words in between which there can be gaps, or wildcards, which match any contiguous sequence of non-whitespace characters. Associated with the gap is a width, which determines how many words the gap can match. An example of a gappy pattern with a gap width of one extracted for a three utterances is shown in Figure 3.

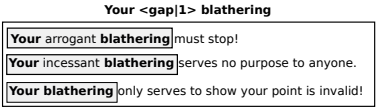Figure 3: An example of a gappy pattern consisting of a single gap which can capture zero or one words and three utterances for which the pattern matches.

The generative model we employ is a modified version of the model introduced by Gimpel and Smith (2011) for machine translation. The main difference is that our model is supervised whereas Gimpel and Smith's model is unsupervised. The supervision in our

model helps to guide the generative process in discovering gappy patterns related to a specific social act.

We assume that we are given a set of utterances $u_{1:k}$ where each utterance is made up of n-words, $w_{1:n}$, and a set of labels $l_{1:k}$ such that label l_i is associated with utterance u_i. Following the terminology of Gimpel and Smith (2011), gappy patterns are represented over the words in the utterances as a color where each word in the utterance has a color assignment, i.e. there is a vector of color assignments $c_{i:n}$ for each utterance. A color C_j is a set made up of the c_i color-word associations such that $C_j = i : c_i = j$. A pattern is built from each color C_j by concatenating the words assigned to the color from left to write inserting gaps between non-adjacent words.

The generative story for a single utterance entails sampling the following:

1. The number of words, n, in an utterance as a $Poisson$ distribution with parameter β.

2. The number of unique colors in an utterance given a $Uniform$ distribution.

3. The color c_i for each word w_i in the utterance as a $Uniform$ distribution.

4. The probability of the pattern associated with each color C_j for utterances with label l as a $Multinominal$ with parameter μ.

Thus, to generate patterns we must calculate the probability of generating an utterance of length n, with m colors, label l, and color assignments $c_{1:n}$ as:

$$p(w_{1:n}, c_{1:n}, m | \beta, \mu) = \frac{1}{Z} (\frac{\beta^n}{n!} e^{-\beta})(\frac{1}{n})(\frac{1}{m})^n \prod_{j=1}^{m} p_\mu(\pi(C_j) \mid l)$$

where Z is a normalization factor.

The multinomial distribution, p_μ, is modeled using a Dirichlet process. A Dirichlet process can be treated as a probability distribution over random distributions which facilitates an unbounded set of parameters $\mu \sim DP(\alpha, P_0)$, where α is the concentration parameter and P_0 is the base distribution. The base distribution is made up of: (1) a $Poisson$ distribution with parameter υ over the number of words in the utterance; (2) a uniform distribution for each word; (3) a uniform distribution over the number of gaps given the number of words; and (4) a uniform distribution over the arrangement of gaps given the number of gaps and words.

Gibbs sampling is used to sample the posterior distribution $p(\{c^{(i)}, m^{(i)}\}_{i=1}^{U} \mid \{w^{(i)}\}_{i=1}^{U}, \upsilon, \alpha)$, where U is the total number of utterances. Gibbs sampling is a form of Markov Chain Monte Carlo (MCMC), which is used to obtain a sequence of random samples from a joint probability distribution of two or more random variables. In particular, Gibbs sampling is used when direct sampling is prohibitive. The Gibbs sampler makes repeated iterations. During each of the iterations it samples a new color for each of the color assignments (c_i). A new color is assigned to c_i by first removing the current color and then choosing from either one the other m colors in the utterance or a creating a new color. The probability of choosing a new color is proportional to:

$$\frac{N_{\pi(\{i\})} + \alpha P_0(\pi(\{i\}) \mid l)}{N + \alpha}$$

where $N_{\pi(\{i\})}$ is the count of pattern π over all utterances with label l and N is the total count of all the patterns. The probability of assigning an existing color j to c_i is proportional to:

$$\frac{N_{\pi(C_j \cup \{i\})} + \alpha P_0(\pi(C_j \cup \{i\}) \mid l)}{N_{\pi(C_j)} + \alpha P_0(\pi(C_j) \mid l)}$$

where $C_j \cup \{i\}$ states that c_i is being added to C_j.

After discovering gappy patterns using the described generative model, we build a binary classifier for each of the social acts we wish to identify. In particular, we use a logistic regression classifier to the learn a set of weights for each of the gappy patterns, which denotes the discriminatory ability of the pattern in identifying the social act. A social act is manifested in an utterance when $H(z) = 1$, where $H(z)$ is calculated as:

$$H(z) = \begin{cases} 1, & \frac{1}{1+e^{-z}} > 0.5 \\ 0, & \frac{1}{1+e^{-z}} \leq 0.5 \end{cases}$$

where $z = \sum_{i=1}^{m} w_i * \phi(\pi_i, w_{i:n})$ and ϕ returns 1 if pattern π_i is present in the utterance made up of words $w_{1:n}$. An utterance is then assigned all social acts whose accompanying classifier results in an $H(z) = 1$.

5 Data Collection & Annotation

We constructed a corpus of 215 social dialogues extracted from English Wikipedia talk pages, public forums, and chat transcripts. A total of 21,067 utterances were extracted from the social dialogues. On average each utterance contained 18.7 words.

A web-based interface was constructed for annotation. The interface listed for a single social dialogue all of the utterances in the order in which they appeared in the dialogue along with speaker information. Social acts were annotated by through the use a drop-down list and allowed for an arbitrary number of social acts to be assigned to an utterance.

Each utterance was annotated by two annotators, who were trained linguists, as either being a manifestation or not of one or more of the eleven social acts described in section 3. In total, 8,149 (38.7%) of the total utterances had at least one of the eleven social acts annotated. On average each utterance was assigned with 0.98 social acts.

We first looked at the inter-annotator agreement for if an utterance exhibited any social act or not. The results are listed in Table 1. The micro-averaged mutual F-Measure was 94.0% which broke down as 84.0% F-Measure for "exhibited" and a 94.0% F-Measure for "Not Exhibited." The results show that expert annotators can reliably determine the presence or absence of social actions in utterances.

Next, we examined the inter-annotator agreement rate for each of the individual social acts. Table 2 shows the number of utterances annotated for each social act, the Cohen's kappa (Cohen, 1960), and the mutual F-measure.

As seen in table 2, the kappa values range from 0.13 to 0.53. In contrast, the kappa values for dialogue acts have been reported as high as 0.76 for ANSWER and as low as 0.15 for COMMITTING-SPEAKER-FUTURE-ACTION (Allen and Core, 1997). More recent work in dialogue act annotation has been performed by Geertzen and Bunt (2006) who report kappa

	F-Measure
Exhibited	84.0%
Not Exhibited	94.0%
micro-Averaged	94.0%
macro-Averaged	89.0%

Table 1: The mutual F-Measure for an utterance exhibiting or not exhibiting any social act.

	# Annotated	Kappa	F-Measure
Agreement	295	0.38	0.76
Challenge Credibility	1,113	0.36	0.33
Disagreement	434	0.46	0.71
Disrespect	367	0.24	0.54
Establish Credibility	364	0.53	0.45
Managerial Influence	2,486	0.23	0.16
Mediation	167	0.26	0.52
Relationship Conflict	399	0.13	0.21
Solidarity	100	0.52	0.42
Supportive Behavior	269	0.36	0.68
Task Conflict	2,802	0.35	0.31

Table 2: The number of annotations and kappa per social act.

values between 0.21 and 1.0 for the top level of a hierarchical dialogue act scheme (Bunt et al., 2010). However, they only calculated kappa for utterances in which both annotators had assigned a dialogue act, i.e. utterances where only one annotator assigned a dialogue act were ignored. In contrast, we calculated our kappa values for all utterances where at least one annotator assigned a social act.

Other work in social acts have seen kappa values in a similar range, such as the Bender et al. (2011) who report kappa values from 0.13 to 0.63. Given the complexities presented by annotating the social intentions of dialogue participants, we believe that the kappa values reported here are acceptable given the accompanying F-Measures.

6 Experimental Results

The utterances in the corpus were labeled with a social act if it was assigned by either of the annotators. The reason for this was two-fold: (1) The definitions of the social acts can be interpreted differently depending on internal thresholding (e.g. how hostile does a remark need to be in order to be classified as a Relationship Conflict?) and no interpretation is truly incorrect. (2) Given the sparsity of annotation for some of the social acts (e.g. Solidarity only having 100 total annotations) made it necessary to include any instance. Experiments were then performed using a standard 80/10/10 split where 80% of the data was used for training, 10% for development, and 10% for testing. We examined an n-gram based approach for comparison to the gappy patterns.

For the gappy pattern method, we constrained the gap width to 2, meaning that the maximum number of words a gap can consume was two. In addition, the sampling process was ran for 1,000 iterations with a burn-in of 50 iterations. A minimum probability of

70% was needed to initiate a new color and a minimum probability of 40% was needed to propagate a color. These parameters were tuned using the development set.

The N-gram based method is a simplification of the gappy pattern method, i.e. it is patterns without gaps. First, n-grams were extracted from the annotated data. Second, the n-grams were pruned using information gain where the exact number of features retained was determined using the development set. Finally, the remaining n-grams were used as binary features in an Support Vector Machine (SVM) classifier with a linear kernel. SVMs are frequently used in text classification and have been shown to give good results for dialogue acts (Hu et al., 2009). We examined using unigrams (1-gram) and unigrams and bigrams (1+2-grams). Other size bigrams were tried as well as incorporating part-of-speech, but resulted in extremely poor performance.

	Precision	Recall	F-Measure
Gappy	49.5%	51.3%	50.4%
1-gram	43.8%	43.2%	43.5%
1+2-grams	38.4%	37.8%	38.1%

Table 3: Micro averaged precision, recall, and F-measure for identifying the 11 social acts using an 80/20 split.

Table 3 lists the micro-averaged precision, recall, and F-measure for the identification of the 11 social acts. As can be seen in the table, the gappy-pattern based approach had an increase of 6.9% in F-Measure over the best n-gram approach (1-gram). This included an increase in precision of 5.7% and in recall of 8.1%. The unigram and bigram (1+2-grams) based method performed the worst, mostly due to the size of data. In contrast, the gappy pattern approach was able to learn a mix of patterns of varying lengths and gaps (up to 2) that were able to separate the social acts. Figure 4 shows example patterns for each of the social acts that were discovered by the generative model and then weighted as a positive indicator by logistic regression.

Conclusion

In this work we have begun a process of revisiting classical theories of dialogue coherence and understanding. Our initial efforts have been to show how the social intentions behind dialogue segments can be understood. Understanding the social goals of individuals is critical to properly parsing dialogue in a modern age of social media. We introduced a set of social acts designed to capture the social intentions of dialogue participants. Our results show that social acts can be reliably understood by annotators and that a novel method for detecting speech acts, based on a generative method for identifying gappy patterns, can achieve results consistent with work in recognizing dialogue acts.

This work creates a foundation for building models of the higher level intentional structure of social dialogue and attention. The intentional structure around the social acts could provide valuable insight into the social goals of a dialogue and the dialogue participants. It is also critical that future work addresses the way in which social attention is modulated within a coherent dialogue.

Agreement	Supportive Behavior	Challenge Credibility
exactly	support	tring bullshit
maybe	works beautifully	been spite
certainly	woman support	reverting move
fair	appreciate	fixing
Disagreement	**Disrespect**	**Relationship Conflict**
violate	simply cowardice	insult <gap> against
your <gap> unacceptable	fooling <gap> fallacy	simple
misinformative	hahaha	amazingly <gap> decision
not agree	quaking	arrogant blathering
Mediation	**Task Conflict**	**Establish Credibility**
insulting each	threshold	myself
false accusations	article title	difference between
former <gap> violates	sources	lead <gap> reflect
should resolve	rebutting reliable	giving <gap> revert
Managerial Influence	**Solidarity**	
discussion	definitely improvement	
also <gap> problematic	good point	
information should	improving thanks	
contentious <gap> problem	reliable	

Figure 4: Example patterns for each social act that were discovered using the generative model and then determined to be a positive indicator by logistic regression.

Acknowledgment

References

Allen, J. and Core, M. (1997). Draft of DAMSL: Dialog Act Markup in Several Layers.

Anderson, C., John, O. P., Keltner, D., and Kring, A. M. (2001). Who attains social status? Effects of personality and physical attractiveness in social groups. *Journal of Personality and Social Psychology*, 81(1):116–132.

Austin, J. L. (1962). *How to do things with words*, volume 7 of *The William James lectures*. Harvard University Press.

Barzilay, R. and Lapata, M. (2005). Modeling local coherence: an entity-based approach. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, ACL '05, pages 141–148, Stroudsburg, PA, USA. Association for Computational Linguistics.

Bender, E. M., Morgan, J. T., Oxley, M., Zachry, M., Hutchinson, B., Marin, A., Zhang, B., and Ostendorf, M. (2011). Annotating social acts: authority claims and alignment moves in wikipedia talk pages. In *Proceedings of the Workshop on Languages in Social Media*, LSM '11, pages 48–57, Stroudsburg, PA, USA. Association for Computational Linguistics.

Blei, D. M., Ng, A., and Jordan, M. (2003). Latent dirichlet allocation. *JMLR*, 3:993–1022.

Bracewell, D., Tomlinson, M., Brunson, M., Plymale, J., Bracewell, J., and Boerger, D. (2012). Annotation of adversarial and collegial social actions in discourse. In *Proceedings*

of the Sixth Linguistic Annotation Workshop, pages 184–192, Jeju, Republic of Korea. Association for Computational Linguistics.

Bracewell, D. B., Tomlinson, M., Shi, Y., Bensley, J., and Draper, M. (2011). 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.

Bramsen, P., Escobar-Molano, M., Patel, A., and Alonso, R. (2011). 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, pages 773–782, Stroudsburg, PA, USA. Association for Computational Linguistics.

Brewer, M. B. and Gardner, W. (1996). Who is this "We"? Levels of collective identity and self representations. *Journal of Personality and Social Psychology*, 71(1):83–93.

Bunt, H., Alexandersson, J., Carletta, J., Choe, J.-W., Fang, A. C., Hasida, K., Lee, K., Petukhova, V., Popescu-Belis, A., Romary, L., Soria, C., and Traum, D. R. (2010). Towards an iso standard for dialogue act annotation. In *LREC*.

Byron, D. and Stent, A. (1998). A preliminary model of centering in dialog. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics - Volume 2*, ACL '98, pages 1475–1477, Stroudsburg, PA, USA. Association for Computational Linguistics.

Campion, M., Papper, E., and Medsker, G. (1996). Relations between work team characteristics and effectiveness: A replication and extension. *Personnel psychology*, 49(2):429–452.

Cassell, J., Nakano, Y. I., Bickmore, T. W., Sidner, C. L., and Rich, C. (2001). Non-verbal cues for discourse structure. In *Proceedings of the 39th Annual Meeting on Association for Computational Linguistics*, ACL '01, pages 114–123, Stroudsburg, PA, USA. Association for Computational Linguistics.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20(1):37–46.

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.

Deutsch, M. (2011). Cooperation and competition. *Conflict, Interdependence, and Justice*, pages 23–40.

French, J. R. P. and Raven, B. (1959). The Bases of Social Power. In Cartwright, D., editor, *Studies in Social Power*, volume 35 of *Studies in social power*, chapter 9, pages 150–167. Institute for social research.

Garrod, S. and Pickering, M. J. (2004). Why is conversation so easy? *Trends in Cognitive Sciences*, 8(1):8–11.

Geertzen, J. and Bunt, H. (2006). Measuring annotator agreement in a complex hierarchical dialogue act annotation scheme. In *Proceedings of the 7th SIGdial Workshop on Discourse and Dialogue*, SigDIAL '06, pages 126–133, Stroudsburg, PA, USA. Association for Computational Linguistics.

Gimpel, K. and Smith, N. A. (2011). Generative Models of Monolingual and Bilingual Gappy Patterns. *Proceedings of the Sixth Workshop on Statistical Machine Translation*, (2009):512–522.

Grosz, B. J. (1978). *Understanding Spoken Language*, chapter Discourse Analysis. Elsevier Science.

Grosz, B. J. and Sidner, C. L. (1986). Attention, Intention, and the Structure of Discourse. *Computational Linguistics*, 12(3):175–204.

Hobbs, J. R. (1979). Coherence and coreference *. *Cognitive Science*, 3(1):90–67.

Hobbs, J. R. (1985). On the Coherence and Structure of Discourse. In *CSLI 85-37*.

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.

Jehn, K. A. and Mannix, E. A. (2001). The Dynamic Nature of Conflict: A Longitudinal Study of Intragroup Conflict and Group Performance. *Academy of Management Journal*, 44(2):238.

Jekat, S., Klein, R., Maier, E., Maleck, I., Mast, M., Berlin, T., Quantz, J. J., and Quantz, J. J. (1995). Dialogue acts in verbmobil. Technical report.

Kamp, H. (1984). A theory of truth and semantic representation. In Groenendijk, J., Janssen, T. M. V., and Stokhof, M., editors, *Truth, Interpretation and Information: Selected Papers from the Third Amsterdam Colloquium*, pages 1–41. Foris Publications, Dordrecht.

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.

Kim, J. and Galstyan, A. (2010). Towards modeling social and content dynamics in discussion forums. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Linguistics in a World of Social Media*, WSA '10, pages 13–14, Stroudsburg, PA, USA. Association for Computational Linguistics.

Mann, W. C. and Thompson, S. A. (1988). Rhetorical Structure Theory: Toward a Functional Theory of Text Organization. *Text*, 8(3):243–281.

Marcu, D. and Echihabi, A. (2002). An unsupervised approach to recognizing discourse relations. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, ACL '02, pages 368–375, Stroudsburg, PA, USA. Association for Computational Linguistics.

Mayfield, E. and Rose, C. P. (2011). Recognizing Authority in Dialogue with an Integer Linear Programming Constrained Model. In *Computational Linguistics*, pages 1018–1026. Association for Computational Linguistics.

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

Petukhova, V. and Bunt, H. (2011). Incremental dialogue act understanding. In *Proceedings of the Ninth International Conference on Computational Semantics*, IWCS '11, pages 235–244, Stroudsburg, PA, USA. Association for Computational Linguistics.

Searle, J. R. (1969). *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press.

Shriberg, E., Dhillon, R., Bhagat, S., Ang, J., and Carvey, H. (2004). The ICSI Meeting Recorder Dialog Act (MRDA) Corpus. In Strube, M. and Sidner, C., editors, *Proceedings of the 5th SIGdial Workshop on Discourse and Dialogue*, pages 97–100, Cambridge, Massachusetts, USA. Association for Computational Linguistics.

Smith, P. and Galinsky, A. (2010). The nonconscious nature of power: Cues and consequences. *Social and Personality Psychology Compass*, 4(10):918–938.

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

Tomlinson, M., Bracewell, D. B., Draper, M., Almissour, Z., Shi, Y., and Bensley, J. (2012). Pursing power in arabic on-line discussion forums. In *Proceedings of the Eighth Conference on International Language Resources and Evaluation*.

Traum, D. R. and Hinkelman, E. A. (1992). Conversation acts in task-oriented spoken dialogue. *Computational Intelligence*, 8(3):575–599.

Webb, N. and Ferguson, M. (2010). Automatic extraction of cue phrases for cross-corpus dialogue act classification. In *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*, COLING '10, pages 1310–1317, Stroudsburg, PA, USA. Association for Computational Linguistics.