

Actions Speak as Loud as Words: Predicting Relationships from Social Behavior Data

Sibel Adali
Rensselaer Polytechnic
Institute
110 8th Street
Troy, New York
sibel@cs.rpi.edu

Fred Sisenda
Rensselaer Polytechnic
Institute
110 8th Street
Troy, New York
sisenf@cs.rpi.edu

Malik Magdon-Ismael
Rensselaer Polytechnic
Institute
110 8th Street
Troy, New York
magdon@cs.rpi.edu

ABSTRACT

In recent years, new studies concentrating on analyzing user personality and finding credible content in social media have become quite popular. Most such work augments features from textual content with features representing the user's social ties and the tie strength. Social ties are crucial in understanding the network the people are a part of. However, textual content is extremely useful in understanding topics discussed and the personality of the individual. We bring a new dimension to this type of analysis with methods to compute the type of ties individuals have and the strength of the ties in each dimension. We present a new genre of behavioral features that are able to capture the "function" of a specific relationship without the help of textual features. Our novel features are based on the statistical properties of communication patterns between individuals such as reciprocity, assortativity, attention and latency. We introduce a new methodology for determining how such features can be compared to textual features, and show, using Twitter data, that our features can be used to capture contextual information present in textual features very accurately. Conversely, we also demonstrate how textual features can be used to determine social attributes related to an individual.

Categories and Subject Descriptors

H.3.1 [Information Storage and Retrieval]: Content Analysis and Indexing; H.5.3 [Group and Organization Interfaces]: Web-based interaction; J.4 [Social and Behavioral Sciences]

General Terms

Algorithms, Measurement, Human Factors

Keywords

Social ties, Social signals, Behavior analysis, Social networks

1. INTRODUCTION

The main problem that we are attempting to address in this paper is the following: given a pair of individuals who are engaged in public discourse using social media, what is the

nature of the relationship between them? On the level of interpersonal relationships, Uzzi classifies ties as arm's-length or embedded [28]. Arm's-length ties function without any prolonged social contact between parties. These ties are an excellent means of cheaply acquiring public information as they do not require any extended communication or social contracts. In embedded ties, information exchanges between actors are dependent on social attachments that produce expectations of trust and reciprocity. These high levels of trust and reciprocity facilitate the sharing of private information. A great deal of work that concentrates on information cascade processes, community creation and evolution tend to concentrate on all links in the network as the same. However, arm's length and embedded links provide different advantages to the individual. The arm's-length ties permit the person to search the network for public information efficiently, while the embedded ties allow the person to share risks and collaborate efficiently with others. Research shows that an individual with both types of links can out-perform others in the network as they can find and integrate diverse public and private information [28]. As a result, algorithms concentrating on prominence of individuals need to concentrate on not only the structure of a network, but also the individual's immediate relationships. As information traveling in more trusted networks is more likely to be believed, information cascade models need to take into account the type of link between individuals. For example, we need to compare cascade processes that start from high trust and embedded links versus arm's length links. Information credibility models [7] that treat follower relationships as vote of confidence by the network can be further refined by distinguishing social ties from those based on reputation. Most work on studying spam [5] [29] tends to disregard the distinction between such processes. Similarly, algorithms that predict trust through trust propagation must concentrate on the nature of the tie as trust will likely propagate differently for different ties [10].

We approach this classification problem with two different toolsets. First, we borrow linguistic analysis tools from social psychology work. Pennebaker [8] [23] shows that the choice of words are strong indicators of the personal and social processes that individuals are engaged in. As a result, one can use the types of words used in exchanges between pairs of individuals as indicators of their relationship and develop classifiers from these categories to social relationships. However, text analysis is in general expensive and fairly sophisticated tools are needed for this type of analysis

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.

WWW 2012, April 16–20, 2012, Lyon, France.
ACM 978-1-4503-1229-5/12/04.

Word usage ¹	<i>words per sentence, words with more than six letters, dictionary words, swear words</i>
Pronouns ¹	<i>e.g. I, you, she, me, mine, themselves, etc.</i>
Verbs ¹	<i>verbs, auxiliary verbs e.g. am, will, have</i>
Tenses ¹	<i>past, present, future</i>
Other function words ¹	<i>articles, adverb, prepositions, conjunctions, negations, quantifiers, numbers</i>
Social Processes ²	<i>family, friend, humans e.g. daughter, buddy, baby</i>
Affective Processes ²	<i>positive & negative emotions, e.g. anxiety, anger, sadness</i>
Cognitive Processes ²	<i>insight, causation, discrepancy, tentative, certainty, inhibition, inclusive, exclusive</i>
Perceptual Processes ²	<i>see, hear, feel</i>
Biological Processes ²	<i>body, health, sexual, ingestion</i>
Relativity ²	<i>motion, space, time</i>
Work ³	<i>job, xerox, majors</i>
Achievement ³	<i>earn, hero, win</i>
Leisure ³	<i>cook, chat, movie</i>
Home ³	<i>apartment, kitchen, family</i>
Money ³	<i>audit, cash, owe</i>
Religion ³	<i>alter, church, mosque</i>
Death ³	<i>bury, coffin, kill</i>
Spoken Categories	<i>assent, nonfluencies e.g. hmm, fillers e.g. blah</i>
Punctuation	<i>.,;:</i>

Figure 1: The 20 textual Categories used in our experiments. Categories marked with ¹ are linguistic processes, ² are psychological processes and ³ are personal concerns. For each category, the features contained in that category or examples are shown.

in each different language. Also, text may not be available in some cases for analysis due to the nature of the medium or as a result of the privacy concerns.

To answer this second challenge, we develop a second toolset based almost social attributes and the statistical analysis of the behavior between individuals. No textual analysis is used for these features. In fact, all the behavioral features rely solely on the timing between the interactions between pairs of individuals. Our methods first parse these interactions into different types of social actions using novel algorithmic methods. We then construct novel behavior measures for each social action based on basic concepts from social psychology. An important factor to consider for these measures is that they are relative to the individual’s social circle and overall behavior. We consider measures based on user status, balance (or assortativity), relative bandwidth/attention dedicated to another, reciprocity of different types of actions. We also add to these measures based on priority and delay between communications that are indicative of the roles the individuals play in a relationship. A large number of these features are novel and to our knowledge, no study has considered these types of features in studying social relationships. We add to these a set of frequently studied social features for our analysis.

Given these two toolsets for studying relationships, we develop a novel methodology for comparing these two: can one set of features be used to predict the other set of features? How can we determine classes of relationships from these two toolsets? Using real data from Twitter, we show that we can use these two toolsets in two different directions to determine the relationships between people very accurately. In fact, behavioral features can be successfully used to predict linguistic categories. Furthermore, the classes identified using behavioral and linguistic features closely resemble each other. As an added bonus, we show that the classes identified by our methods closely resemble the two that have

been studied at great length in the literature: embedded vs. arm’s length [28].

The rest of the paper builds a case for the following take away message: If the statistical nature of the interaction between *A* and *B* is reciprocal with *A* paying attention to *B* and prioritizing him over others (as measured by our behavioral features), then it is likely that the relationship is of a personal nature (as measured by linguistic analysis of the communications). On the other hand, if the interaction is non-reciprocal but between individuals of equal social stature (again as measured by our behavioral features), then it is likely that the relationship is of a formal nature (for example information exchange, work related, etc.).

2. LINGUISTIC FEATURES

To study the relationships between individuals, first we use a set of features based on linguistic characteristics of text, using a tool called LIWC [27] (Linguistic Inquiry and Word Count). LIWC provides a number of linguistic categories and classifies words in English into these categories. Various empirical studies have shown that the categories in LIWC are effective in determining differences between individuals, social relationships, especially attentional focus and emotionality of the relationship. Because of this, LIWC features are a good match for the study provided in this paper. These features were used for detecting personality types in earlier studies [12] [13] and were shown to be among the most useful features for this purpose.

There are about 83 individual categories provided by LIWC based on a dictionary composed of more than 3000 words and word stems. Each word or word-stem defines one or more word categories or subdictionaries. For example, the word “cried” is part of four word categories: sadness, negative emotion, overall affect, and a past tense verb. Hence, if it is found in the target text, each of these four subdictionary scale scores will be incremented. As a result, the features

from the same category tree are correlated. We eliminated word count, filename and segment features from this list as they do not provide any specific linguistic context. This left us with 80 LIWC features some of which are top level categories.

We also consider a set of 20 categories that these 80 features fall into. Our categories are mostly second level categories from LIWC. However, we decided to break down the large category in linguistic processes to three main groups. We also grouped all punctuation into a single category. The resulting categories are all disjoint, they do not contain any repetitive counts due to hierarchical relationships. These categories and a list of their subcategories are shown in Figure 1.

3. SOCIAL AND BEHAVIORAL FEATURES

To understand the social relationship between pairs of individuals, we examine their social network. Most work on analyzing social media consider whether the individuals are friends or followers of each other, and use simple behavioral measures such as the number of messages between them. Measures that indicate how popular users are such as the total number of friends and followers are also used. These measures while useful are also quite noisy. A pair can be friends with each other but rarely exchange messages. A person may send a lot messages, but never receive a response. A person with many friends may not forward any of their messages. Furthermore, the existence of broadcast messages alongside with direct messages in Twitter provides an additional complication to this type of analysis: can we extract useful friendship information from broadcast messages? In this paper, we implement a set of social and behavioral features, some are based on known social measures and some are a completely new set of features based on how people act towards each other.

To understand the social relationship between pairs of individuals in social media, we first look at the interaction between them. We consider a trace of interactions of the form: $\text{Trace}(\text{sender}, \text{receiver}, \text{time}, \text{type})$ from a set of individuals. Each tuple in Trace is of the form: (A, B, t, z) which means that A sent a message to B at time t of type z . In the case of Twitter, messages can be either directed or broadcast. In emails, all messages are directed. Next, we extract for each pair (A, B) of users, a set of social actions. These are the actions that individuals participate in the given social media environment. In the case of Twitter, we consider three main classes actions:

- PAIR is just any exchange between the two individuals. It considers directed messages.
- CONV considers “conversations”, i.e. sustained exchange of directed messages between the two individuals in a short amount of time.
- PROP considers “propagations”, i.e. messages of any kind from A to B that are later propagated by B to someone else.

As trace does not contain message content, we instead develop algorithmic methods to compute the contents of the CONV and PROP actions for the given pair. We first describe these below.

3.1 Parsing trace into social actions

Suppose we are analyzing the relationship between two pairs of individuals (A, B) by analyzing contents of Trace . Suppose $\text{Msg}_d(A, B)$ is the list of all tuples in Trace of the form $(A, B, -, \text{directed})$ where $-$ stands for don’t care. We also use $\text{Msg}(A, B)$ to denote the tuples of the form $(A, B, -, -)$ containing direct and broadcast messages, $\text{In}(A)$ to denote all messages to A from anyone (of the form $(-, A, -, -)$) and $\text{Out}(A)$ to denote all messages from A to anyone (of the form $(A, -, -, -)$).

We look at three types of behavior:

PAIR: Pair behavior Given a pair (A, B) , the Pair set is given by $\text{PAIR}(A, B) = \text{Msg}_d(A, B) \cup \text{Msg}_d(B, A)$. In other words, the set contains all directed messages between the two users.

CONV: Conversation behavior For conversations, we group the tuples in $\text{PAIR}(A, B)$ into conversations,

$$\text{CONV}(A, B) = \{C_1, \dots, C_n\}$$

where each conversation C_i is a set of consecutive messages from $\text{PAIR}(A, B)$ and $C_i \cap C_j = \emptyset$ for all $i \neq j$. We consider all conversations of size at least 2 in our experiments.

To find a conversation, first we find the average time interval between any two consecutive messages in $\text{PAIR}(A, B)$ called τ . Then, any pair of messages that occur close to each other (time interval between the pair is less than $c * \tau$ for some smoothing factor c) are considered part of the same conversation. We use $c = 3$ in our tests. Only users with at least one conversation will have conversation related features. The complexity of computing this feature is $O(|D| \log |D|)$ where $|D|$ is the size of the trace stream. All features based on conversations are computed as a function of the conversation construction and hence do not incur additional costs.

PROP: Propagation behavior Given a pair of users (A, B) , the set $\text{PROP}(A, B)$ contains pairs of the form (t_1, t_2) such that $t_1 \in \text{Msg}(A, B)$ and $t_2 \in \text{Out}(B)$ and t_2 is most likely a propagation of t_1 .

To find potential propagations by B , we use a linear time maximum matching algorithm developed in [4] from $\text{In}(B)$ to $\text{Out}(B)$ satisfying a causality constraint with respect to time: t_2 should come after t_1 . Given the output of the match, we compute the subset of messages which were from A and were sent to someone other than A . These pairs of messages are the messages in $\text{PROP}(A, B)$. Note that we do not actually check the message to see if it was a retweet to verify that it was a propagation. In previous work, we have shown that our notion of propagation is correlated with true propagation in the form of retweets. However, even in the absence of retweeting, our match also implies a correlation between messages received by B from A and messages sent by B .

3.2 Behavioral Features

Given these different sets $\text{PAIR}, \text{CONV}, \text{PROP}$ of social actions, we construct a set of features based on behavior. The following are the main classes of features for a pair (A, B) that measure the relationship A has with B . Note the directionality of this measurement. A specific type of (A, B) relationship is not guaranteed to be symmetric to the relationship (B, A) (for example A may not know B well, but B may follow A closely).

USER features		Relative attention features	
U-Follow	<i># users who follow the user (log)</i>	A-ATTN-Conv	<i># messages in Conv(A, B)</i>
U-Friend	<i># friends of the user (log)</i>	A-ATTN-#Conv	<i># conversations in Conv(A, B)</i>
U-Years	<i># years on Twitter</i>	A-ATTN-CMsg	<i>avg. # msg per conv. in Conv(A, B)</i>
U-Out	<i># messages sent (log)</i>	A-ATTN-CMax/CMin/CStd:	<i>max/min/std of # msgs in conv.</i>
U-Fav	<i>ratio of # messages favorited to U-Out</i>	A-ATTN-Trust	<i>conversation trust</i>
U-Spacing	<i>mean # hours between messages</i>	A-ATTN-Msg	<i>Msg(A, B)</i>
U-Retweet	<i>ratio of retweet messages</i>	A-ATTN-INB	<i>In(B)</i>
U-Directed	<i>ratio of directed messages</i>	A-ATTN-OutA	<i>Out(A)</i>
U-URL	<i>ratio of messages with URLs</i>	A-ATTN-Prop	<i>Prop(A, B)</i>
U-HashTag	<i>ratio of messages with hashtags</i>	A-ATTN-PropB	<i>Prop(-, B)</i>
U-Mention	<i>ratio of messages with mentions</i>	B-ATTN-W	<i>% of A's messages B finds worthy to propagate, Prop(A, B)/Out(A) ∩ In(B)</i>
U-TextLength	<i>mean length of messages divided by 140</i>	B-ATTN-E	<i>% of B's propagation energy spent on A, Prop(A, B)/Prop(-, B)</i>
U-PropFrom	<i># people user propagates/U-Follow</i>		
U-PropTo	<i># people propagate from user/U-Friend</i>		
U-Conv	<i># people user converses with/U-Friend</i>		
Other	<i>mean and std of all behavioral features across all people user interacts with</i>		
Balance & reciprocity features		Time, priority & delay features	
BAL	<i>degree similarity of users</i>	TIME-AB	<i>avg. response time from A to B</i>
RECIP-C	<i>reciprocity of conversation messages</i>	TIME-BA	<i>avg. response time from B to A</i>
RECIP-P	<i>reciprocity of Prop(A, B), Prop(B, A)</i>	TIME-tau	<i>avg. response time in conversations</i>
		PRI-AB	<i>how much A prioritizes B over others</i>
		PRI-BA	<i>how much B prioritizes A over others</i>
		DEL-AB	<i>delay in conversation from A to B</i>
		DEL-BA	<i>delay in conversation from B to A</i>

Figure 2: List of statistical behavioral features in various categories with abbreviations and short explanations

- **USER: user's network (USER-A, USER-B)** measures a number of social and behavioral features of the user, not specific to a pair. Examples of these are the number of friends (those the user follows), followers (people who follow the user), total number of days they have been using Twitter, etc. We implement the features from related work [26] and add others to the list.
- **A-ATTN: A's relative attention** measures how much B is getting A 's total attention. This measure is computed by the number of A 's messages satisfying a condition, and normalized by the total messages sent by A . Relative attention is a well known measure in social science [31]. We consider the relative attention for a number of different actions.
- **B-ATTN: B's relative attention** measures how much of B 's attention is given to messages from A . This measure is computed by the number of B 's propagations of A 's messages, and normalized by the total number of messages from B .
- **BAL: Balance (assortativity)** measures the degree similarity of two users. Newman [21] shows that nodes in many different networks tend to be connected to nodes that are alike each other in terms of degree similarity.
- **RECIP: Reciprocity** measures to which degree a node reciprocates the actions of another. Reciprocity is studied in social science literature a great deal [20] [22] as a foundational principle for the formation and maintenance of ties [30]. Embedded and reciprocal relationships [15] form a type of social capital between individuals. Given this capital, they can exchange privileged information and trust each other to accomplish tasks. Both balance and reciprocity are measured by entropy as described below.
- **TIME:** measures the actual time in hours it takes for a user to respond to another person. Faster response times indicate higher priority. The time measures are absolute, they are not relative to one's network. As a result, they favor individuals with fast response time overall as well as ties with fast response time. Overall, fast response time can be considered as an indicator of the other person's reliability in the relationship.
- **PRI: Priority** measures to which degree a person prioritizes another person over all their acquaintances. It shows preference among the ties. Priority is computed by inversions: suppose C sends a message to B and later A sends a message to B . If B responds to A before responding to C , then B prioritizes A over C by one message. Hence, B has inverted the reply ordering from CA to AC . Priority is given by the total number of inversions (i.e. messages) received by a user A from user B , divided by $Out(B)$. Priority together with reciprocity is a strong indication of tie strength between two individuals.
- **DEL: Delay** measures how much a user is typically delayed to get an answer or how many other messages are prioritized over a message from the given user. Delay measures the opposite of priority: how much a person should wait for the other. Delay from A to B is measured in terms of the ratio between number of messages A sends out after receiving a message from B and replying back to B , divided by $Out(A)$. Note that arm's length relationships [28] tend to be non-reciprocal and persist despite long delays. However,

reciprocal ties need short delays to ensure their persistence. We compute priority and delay for conversation actions. Timing, priority and delay [16] are frequently used signals in analyzing relationships in anthropology.

Balance and reciprocity measures both compare two values x_1, x_2 . These two values are reciprocal or balanced when $x_1 = x_2$. When computing balance, we use degree for x_1 and x_2 , i.e. number of followers of individuals. To compute reciprocity, we set x_1 and x_2 to be measures of some action: x_1 from A to B and x_2 the same action from B to A . For example, reciprocity of message propagation. We use the entropy value for both measures as introduced in our previous work [1] and later used in related work on reciprocity [30] as follows:

Given a value x_1 compared to x_2 , we compute balance with entropy $H(x_1 : x_2) = -p \log p - (1 - p) \log(1 - p)$ where $p = x_1 / (x_1 + x_2)$. Entropy is highest when $x_1 = x_2$, i.e. when the balance is highest, and zero when $x_1 = 1$ and $x_2 = 0$.

A-ATTN-Trust measure combines a number of features into a single measure: longer, higher number and more balanced conversations are signs of a strong link. Hence, $A - \text{Attn} - \text{Trust} = \sum_{C \in \text{Conv}(A,B)} |C| * H(|C \cap \text{Out}(A)| : |C \cap \text{Out}(B)|)$ where the entropy computes the balance of the messages in the conversation from A and the messages from B .

Given these various types of features and various social actions, we compute a set of social and behavioral features as shown in Figure 2. Note that some features like reciprocity, balance and conversation features are symmetric. However, propagation features are asymmetric by definition. For a given pair A, B , we consider features for (A, B) (for A 's relationship to B) and (B, A) (for B 's relationship to A). We also consider a number of user features based on social attributes of the user as well as the average and standard deviation of the behavioral features across all the different people the user interacts with. After adding all these features for A and B , we arrive at a total of 130 features for behavior.

4. ANALYSIS

Our aim in this paper is to discover the main relationship categories between pairs of users. To accomplish this, we take a set of user pairs and collect relevant activity between these users. This constitutes our **Trace** relation. For this set, we construct two matrices:

- M_F is a $z \times n_F$ matrix containing all the social and behavioral features as columns ($n_F = 130$ in our case) and different user pairs as rows.
- M_T is a $z \times n_T$ matrix containing all the text features as columns and different user pairs as row ($n_T = 80$ for most tests, but we also consider only the top 20 categories in some others). If row i in M_F corresponds to pair (A, B) , then the corresponding row in M_T is constructed by processing the messages from A to B including broadcasts.

First, we would like to find for a specific category in M_T , the most relevant features in M_F . To accomplish this, we use Forward subset selection based regression (FSSreg). FSSreg performs a regression using an input matrix X and target vector y and results in weights w of the predictor $x^T w$. The

weights w are obtained using a greedy forward stepwise regression to minimize the leave-one-out cross validation error. Specifically, suppose that k features from X have been selected. We select the next feature to minimize the LOO-CV error assuming the first k features are selected. If the LOO-CV error decreases with the $k+1$ th feature, then the process continues. Otherwise it stops and we output the sparse regression vector w using only the k selected features.

If the input is a matrix Y instead of a vector y , then FSSreg is performed independently on each column of Y to output a set of sparse weights for each column of Y .

We run FSSreg comparing each column of M_T to M_F to find behavioral features that are good predictors of word usage in the text. We call this matrix $TF : n_T \times n_F$ for “text to feature” mapping. We have that $M_F \approx M_T * TF$.

We also run FSSreg in the opposite order, comparing each column of M_F to M_T to find what type of word usage is a good indicator of a given behavioral feature. We call this matrix $FT : n_F \times n_T$ for “text to feature” mapping. We have that $M_T \approx M_F * FT$.

These two analyses give us a mapping between features of M_F and M_T . In other words, we show equivalence between any feature from one feature set and a linear combination of features from the other set. The next question we ask is whether we can find groups of these features that are prominent in the given dataset. As an example, consider TF which contains weights that show the prominence of behavioral features for each text feature in M_T . A positive weight in TF means that a specific textual category is positively correlated with this feature, and a negative weight implies a negative correlation. The absolute magnitude shows the relevance of the figure. Based on TF , we compute distances between any two behavioral features f_i, f_j as follows:

$$\text{dist}(f_i, f_j) = \sum_{1 \leq k \leq n_T} (\text{abs}(TF(k, i) - TF(k, j)))$$

which finds the sum of absolute differences of the weights. The higher the distance, the more different are the weights for these two behavioral features, which means that textual features for them are very different. The question we would like to ask now is whether there are natural categories of behavioral features that indicate similar text usage. In other words, we would like to find a clusters of social and behavioral features that have low distance to each other, but are far away from the other features.

To find this, we compute a graph G_{TF} where each node is a behavioral feature f . We add an edge between two nodes f_i, f_j if the distance $\text{dist}(f_i, f_j) < \epsilon$ for a given threshold ϵ . We run the Fast Community clustering algorithm [9] on this graph to find clusters that satisfy the modularity constraint. The TF clusters contain behavioral features that indicate similar word usage. We analyze those to see if there is a natural clustering of word usage and behavior. For each cluster, we list the behavioral features. We also look at the mean weights of textual features for each cluster to see what type of text is relevant to the cluster.

We also compute the opposite, we analyze the FT matrix which shows for a given behavioral feature, the set of relevant text features. We now cluster text features to find word usage that point to similar social features. We analyze these again to see what types of clusters are produced.

This analysis goes a step further than linear regression com-

monly used in similar types of analysis. We also use the regression results to find groups in the data that are naturally represented by two distinct pairs of features.

5. EXPERIMENTAL SETUP

We use Twitter to test our methods. We take a set of 45 random seed users and we find all their followers and friends. The expanded set contains 15664 users. We then collect all available statuses (i.e. messages) for the expanded user set. We can access only public information for both the status collection and the expansion of the user set. A large majority of our user set had public accounts, so we did not have problems growing the set or collecting statuses. We collect all statuses (up to 3200) from this group of users which resulted in about 21 million statuses for our analysis. From this initial set, we choose pairs (A, B) such that A and B both have at least 5 messages and A sent at least one directed message to B . This ensures us that A and B have significant presence on Twitter and that A knows B . Note that in our set, for a pair (A, B) , it is possible that B never writes to A . This could still be a significant relationship if for example B forwards messages from A all the time. This additional constraint reduced the number of pairs in our experiments to 19,933 pairs.

It is also possible for both pairs (A, B) and (B, A) to be present in our set. But, (A, B) contains A 's messages to B whereas (B, A) contains B 's messages to A . Our analysis tries to examine the word usage of A as a function of the relationship A and B have.

Note that M_T features are normalized to be in the same range. The M_F features are also mostly in the 0-1 range. In fact all measures in terms of the number of messages are scaled to A 's energy, i.e. $OUT(A)$, measures in terms of number of friends and followers that A communicates with are scaled by the number of friends and followers A has. All other features are scaled to the same range except for features in the log scale. We run the FSS algorithm from M_T to M_F to find TF for 80 text features in the original LIWC set. We then use these features to find clusters of the social features, denoted by $SocialClusters(TF)$. We set $\epsilon = 1000$ experimentally, though we also note that the algorithm was robust to various values of ϵ . The algorithm produces two disjoint clusters c_F^1, c_F^2 .

In the opposite direction, we find FT by mapping each feature of M_F to a text features in M_T . For this set of experiments, we only choose the text features corresponding to the top 20 categories shown in Figure 1. We then use this set to compute clusters of text categories, denoted by $TextClusters(FT)$. We find two clusters in the opposite direction as well: c_T^1, c_T^2 .

6. RESULTS

6.1 Identifying behavior corresponding to text

First, we analyze the set of behavioral features against the text features. For each text feature, we find the relevant behavioral features and compute TF . We then cluster the social behavioral features based on the similarity of the most significant text categories. The clusters are given in Figure 3. We note that the clusters break the social behavioral features into two distinct groups. Delay, priority and reciprocity of actions mostly lie in cluster c_F^1 . Also, the user

Cluster c_F^1 (embedded)	
User	U-Friend, U-Out, U-Fav, U-Spacing, U-HashTag, U-Mention, U-URL, U-PropTo, U-PropFrom, U-Conv, and standard deviation of some behavioral features (for both users)
Balance & reciprocity	RECIP-P
A Attention	For A: A-ATTN-Conv, A-ATTN-#Conv, A-ATTN-CAvg, A-ATTN-CMax, A-ATTN-CMin, A-ATTN-Trust, A-ATTN-Prop, For B: A-ATTN-Msg, A-ATTN-Prop, A-ATTN-PropB
Delay	DEL-AB, DEL-BA
Priority	PRI-AB, PRI-BA
Cluster c_F^2 (arm's length)	
User	U-Follow, U-TextLength, U-Years, mean and std of B-ATTN-W, B-ATTN-E, BAL, TIME-AB, PRI-AB (for both users)
A Attention	A-ATTN-INB, A-ATTN-OutA
B Attention	B-ATTN-W, B-ATTN-E (for both users)
Balance	BAL
Reciprocity	RECIP-C
Time	TIME-AB, TIME-BA, TIME-tau

Figure 3: Clusters of behavioral features based on TF , i.e. similarity of word usage

features in this cluster measure the number of friends, the type of messages written and the overall attention A pays to B . We consider this cluster a friendship relation, or an embedded relation. It requires reciprocal activity to sustain. Cluster c_F^1 in contrast is dominated by balance in conversations and degree, the amount attention given by B to A , i.e. A 's reputation. Also, the timing between messages is significant. However, message quantity is not significant. In fact, long delays are possible in this group where despite the lack of reciprocity of propagations and possible lack of large volume of messages, the relationship can be sustained. If we look at the user features, we see features like the number of years on Twitter and the number of followers. Both are indicator of status in Twitter. This coupled with assortativity signals relationships between individuals of similar social status, but not as personal a relationship as those in Cluster c_F^1 . This is typical of an arm's length relationship. Furthermore, other user features belonging to this cluster include standard deviation of timing, priority and attention features. This signals that individuals likely have different types of relationships in their network. The availability of these type of dual types of relationships is a source of network power, i.e. access to different pools of information. Given the clusters match almost perfectly the definition of arm's length and embedded relationships, we now ask: what type of topics are discussed in each group? To find this, we find the average weight of the text categories for each cluster. We consider the top 20 categories as discussed earlier for this. If the computed weight is positive, this means that the feature is positively correlated with the pairs in cluster. Negative weights indicate the reverse relation. We show the

Category	c_F^1 Weight	c_F^2 Weight
punctuation	9.16	-4.17
word usage	2.34	96.59
spoken categories	0.76	0.68
social processes	0.51	-0.23
affective processes	0.38	5.00
religion	0.37	-0.20
biological processes	0.27	0.44
home	0.20	0.49
leisure	0.14	4.41
death	0.05	0.13
pronouns	-0.29	9.11
perceptual processes	-0.44	3.34
achievement	-0.73	7.28
money	-0.75	4.21
relativity	-1.35	16.61
cognitive processes	-1.93	11.88
work	-2.66	15.67
other function words	-2.75	22.75
tenses	-3.14	24.57
verbs	-8.79	66.36

Figure 4: Weights of categories based on $SocialClusters(TF)$, clustering of behavior features from TF .

average weights of the text category features in Figure 4. An interesting perspective is offered by categories that are negatively correlated in one cluster and positive in the other. We exclude two outlier categories: verbs and word usage which are both positively correlated C2 and plot the remaining categories in a 2-d graph in Figure 5. In this figure, top left corner represents the categories significantly in Cluster c_F^2 (blue) and the bottom right corner represents the categories significantly in Cluster c_F^1 (red). We see that Cluster c_F^1 is dominated by topics that represent personal discussion, while Cluster c_F^2 mostly contains categories that represent formal and non-personal discussion. This also confirms our initial assessment that cluster c_F^1 represents embedded ties which enables the nodes to exchange privileged/personal information. On the other hand, cluster c_F^2 represents more impersonal discussion which likely does not require trust.

6.2 Identifying text corresponding to behavior

Now that we have established that we can fairly accurately determine classes of ties by examining text features and the associated behavioral features. We now ask the opposite question. What if the textual features are not known? How well can we determine word usage based solely on the behavior? To accomplish this, we run the FSSreg for each behavioral feature and find the textual categories that are best modeled by that specific feature. Now, given these mappings in FT , we compute a graph of text features and cluster these categories into $TextClusters(FT)$ and find two clusters of text categories c_T^1, c_T^2 . For this experiment, we only use the top 20 categories to suppress the noise due to the hierarchical nature of the text features. The text clusters are given in Figure 6. Cluster c_T^2 obtained from FT is very similar to highly positive text categories corresponding to c_F^1 . Both share the categories: cognitive processes, relativity, verbs and word usage. With the exception of punctuation and pronouns, the break down of these features are

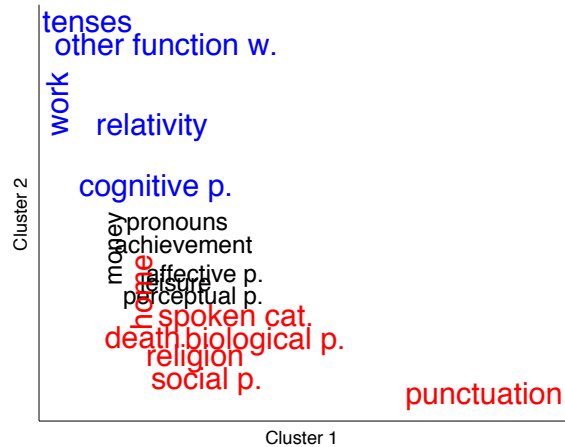


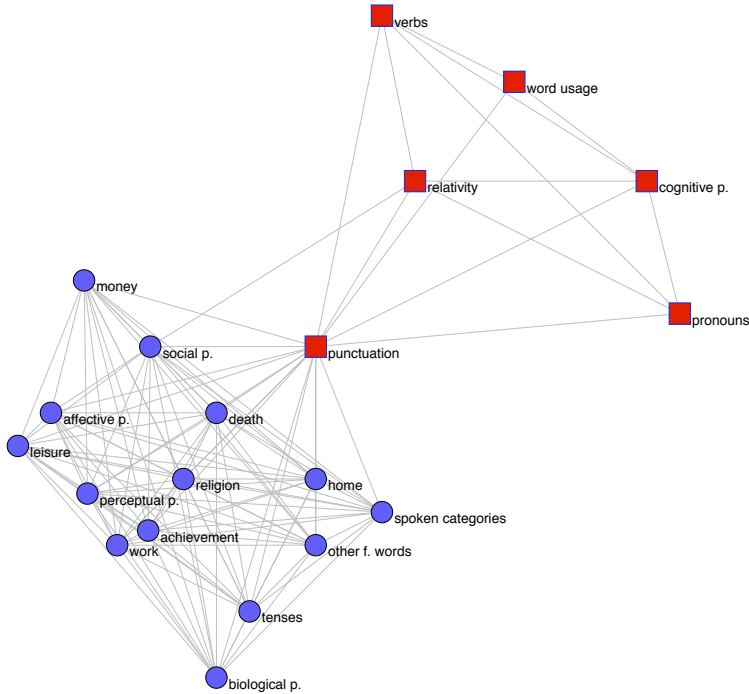
Figure 5: Weights of categories based on $SocialClusters(TF)$. Top left corner (blue) is categories positively correlated to Cluster 2 and negatively correlated to Cluster 1, bottom right corner (red) are categories positively correlated to Cluster 1 and negatively correlated to Cluster 2

	BAL	RECIP	A-Attn	B-Attn	TIME	PRI	DEL
Word Usage		M	+	+	-	-	-
Pronouns	+	-	-	+	-	+	+
Verbs			+				
Tenses	+				+		
Other Function	+		+	-	+	+	+
Social P.	-	+	-	+	+		+
Affective P.		M	-	+	+	-	+
Cognitive P.	+	+	+	-	+		-
Perceptual P.		+			+	+	+
Biological P.	+		M			+	+
Relativity	+		-	+	+	+	+
Work	+	+	+	-	+	-	-
Achievement	-						+
Leisure			-	M		+	+
Home				M		+	
Money	-		-				
Religion	-						
Death							
Spoken Cat.		-	-	+	-	+	-
Punctuation	+	-	+				

Figure 7: The relationship between social and text features based on FT . +/-M indicate a positive/negative/mixed relationship for features in the given class.

almost identical to the clusters determined by text. This is a very striking result. We are able to determine textual categories using behavior alone almost as good as having the text itself. This result leads us to our title: actions speak almost as loud as words.

We also analyze which F features are positively and negatively correlated with each cluster in $TextClusters(FT)$. For each category in that cluster, we find the average weights of the behavioral features from FT . Again, positive (or negative) weights indicate a positive (or negative) correlation. We present a short summary of these in Figure 7 below. First, we note that almost all behavioral features have significant weights in more than half of the categories. Furthermore, the signs of the categories follow the similar trends for TF clusters. For example, PRI and DEL features tend to favor more personal features, while BAL and RECIP favor cognitive features. To illustrate this further, we find the top 100 pairs based on the weights for F for each cluster and construct a word cloud for each cluster as shown in 8. In



Cluster 1		Cluster 2
achievement	other f. words	cognitive p.
affective p.	perceptual p.	punctuation
biological p.	religion	pronouns
death	social p.	relativity
home	spoken cat.	verbs
leisure	tenses	word usage
money	work	

Figure 6: *TextClusters(FT)*: Clusters of textual categories obtained from the *FT* measures, by mapping behavior to text. High distances indicate dissimilar features (very large distances are removed). Cluster 1 (blue circles) corresponds to intimate conversations, cluster 2 (red boxes) corresponds to formal information exchange.

these word clouds, we removed articles and stop words. Furthermore, we removed the common words to show only the words that distinguish the two clusters (for example, if the word think appears 10 times in cluster 1 and 12 times in cluster 2, we remove it from cluster 1 and cluster 2 contains only 2 occurrences.) We can see that there is a significant distinction between the words in each cluster. Words in cluster 1 are emotional and personal, while cluster 2 contains more impersonal and conceptual words. Given that these clusters are found strictly using social and behavioral features, we can claim that the behavioral features are almost equivalent to textual features.

7. RELATED WORK

To our knowledge, no study such as ours exist in the literature. Despite the large body of work analyzing personal relationships in social media and Twitter in particular, behavioral features of the type we propose here have not been studied. The exception to this is our earlier work [1] in which we introduce the parsing methods for conversations and propagations, and introduce three of the features given here (A-ATTN-Trust, B-ATTN-W, B-ATTN-E). While we also introduce entropy, we do not study it. The work in [1] concentrates on the relationship between these features and retweet behavior. This current paper introduces many new features, new analysis methods and poses a very different problem.

Some of the work on analyzing conversational patterns in Twitter either concentrate on the analysis of what the users use Twitter for [6, 19], how textual content of dialogs change

over time [17, 3]. However, this type of work does not try to use the time characteristics to identify the nature of the relationship.

To our knowledge, there is no work on determining the types of relationships. The work on credibility [7] and retweets [26] consider a subset of our user features based on counts and overall status. LIWC features have been used to detect personality in Facebook [13] and Twitter [12]. These two pieces of work and earlier work on predicting social ties from social media [11] based on Facebook mostly look at features like number messages, time since last message, reciprocity of links, number of mutual friends, etc.

The features we introduce here expand on these with measures of balance and reciprocity of a variety of actions. The notion of reciprocity has been studied in the literature a great deal [20, 22, 30] and various measures have been proposed. Our measure tends to work well for the problem studied here. We also introduce a set of new concepts such as the mean time between messages, delay and priority. Furthermore, we normalize our measures for each individual to look at relative attention. This step makes a significant difference in the results. Our methods of normalization are similar to those used by B. Uzzi [31, 24] for extracting real ties from noisy communication data. Furthermore, our statistical methods to parse a communication trace into social actions without considering the textual content of the messages are unique in the literature.

Our tests show that these features provide valuable additional information beyond the simple count or reciprocity of messages. Considering relationships based on the timing

9. REFERENCES

- [1] S. Adali, R. Escrava, M. Goldberg, M. Hayvanovych, M. Magdon-Ismael, B. K. Szymanski, W. A. Wallace, and G. M. Williams. Measuring behavioral trust in social networks. In *Proc. International Conference on Intelligence and Security Informatics (ISI)*, pages 150–152, 2010.
- [2] S. Adali, X. Lu, M. Magdon-Ismael, and J. Purnell. Prominence ranking in graphs with community structure. In *International AAAI Conference on Weblogs and Social Media (ICWSM’11)*, 2011.
- [3] B. D. Alan Ritter, Colin Cherry. Unsupervised modeling of twitter conversations. In *Proceedings of HLT-NAACL*, 2010.
- [4] J. Baumes, M. Golberg, and M. Magdon-Ismael. Efficient identification of overlapping communities. In *Proc. International Conference on Intelligence and Security Informatics (ISI)*, 2006.
- [5] F. Benevenuto, G. Magno, T. Rodrigues, and V. Almeida. Detecting spammers on twitter. In *CEAS 2010 - Seventh annual Collaboration, Electronic messaging, Anti- Abuse and Spam Conference*, 2010.
- [6] D. Boyd, S. Golder, and G. Lotan. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In *HICSS-43. IEEE: Kauai, HI, January 6.*, 2010.
- [7] K. Canini, B. Suh, and P. Pirolli. Finding credible information sources in social networks based on content and social structure. In *2011 IEEE International Conference on Privacy, Security, Risk, and Trust, and IEEE International Conference on Social Computing*, 2011.
- [8] C. Chung and J. Pennebaker. *The psychological function of function words*, pages 343–359. New York: Psychology Press, 2007.
- [9] A. Clauset, M. Newman, and C. Moore. Finding community structure in very large networks. *Phys. Rev. E*, 70(6):066111, 2004.
- [10] T. DuBois, J. Golbeck, and A. Srinivasan. Predicting trust and distrust in social networks. In *IEEE International Conference on Social Computing*, 2011.
- [11] E. Gilbert and K. Karahalios. Predicting tie strength with social media. In *Proceedings of the 27th international conference on Human factors in computing systems*, pages 211–220, 2009.
- [12] J. Golbeck, C. Robles, M. Edmondson, and K. Turner. Predicting personality from twitter. In *2011 IEEE International Conference on Privacy, Security, Risk, and Trust, and IEEE International Conference on Social Computing*, 2011.
- [13] J. Golbeck, C. Robles, and K. Turner. Predicting personality with social media. In *Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems, CHI EA 2011*, pages 253–262, 2011.
- [14] M. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78:1–22, 1973.
- [15] M. Granovetter. Economic action and social structure: The problem of embeddedness. *American Journal of Sociology*, 91:481–510, 1985.
- [16] E. T. Hall. *The Silent Language*. Anchor, 1973.
- [17] C. Honey and S. Herring. Beyond microblogging: Conversation and collaboration via twitter. In *HICSS’09. 42nd Hawaii International Conference on System Sciences*, 2009.
- [18] L. Hong, G. Convertino, and E. Chi. Language matters in twitter: A large scale study. In *International AAAI Conference on Weblogs and Social Media (ICWSM’11)*, 2011.
- [19] A. Java, X. Song, T. Finin, and B. Tseng. Why we twitter: understanding microblogging usage and communities. In *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis*, pages 56–65, 2007.
- [20] D. Lauterbach, H. Truong, T. Shah, and L. Adamic. Surfing a web of trust: Reputation and reciprocity on couchsurfing. In *Proceedings of the 2009 International Conference on Computational Science and Engineering*, volume 4, pages 346–353, 2009.
- [21] M. E. J. Newman. Mixing patterns in networks. *Physical Review E*, 67, 2003.
- [22] E. Ostrom and J. W. (editors). *Trust and Reciprocity*. Russell Sage Foundation, 2005.
- [23] J. Pennebaker. *The Secret Life of Pronouns: What Our Words Say About Us*. Bloomsbury Press, 2011.
- [24] M. Rivera, S. Soderstrom, and B. Uzzi. Dynamics of dyads in social networks: Assortative, relational, and proximity mechanisms. *Annual Review of Sociology*, pages 91–115, 2010.
- [25] S. Saavedraa, K. Hagerty, and B. Uzzi. Synchronicity, instant messaging, and performance among financial traders. *Proceeding of the National Academy of Sciences (PNAS)*, 2011.
- [26] B. Suh, L. Hong, P. Pirolli, and E. Chi. Want to be retweeted? large scale analytics on factors impacting retweet in twitter network. In *2010 IEEE International Conference on Privacy, Security, Risk, and Trust, and IEEE International Conference on Social Computing*, 2010.
- [27] Y. R. Tausczik and J. W. Pennebaker. The psychological meaning of words: Liwc and computerized text analysis methods. *Journal of Language and Social Psychology*, 29(1):24–54, 2010.
- [28] B. Uzzi. Social structure and competition in interfirm networks: The paradox of embeddedness. *Administrative Science Quarterly*, 42:35–67, March 1997.
- [29] A. Wang. Don’t follow me: Spam detection in twitter. In *Intl Conference on Security and Cryptography (SECRYPT)*, 2010.
- [30] C. Wang, A. Strathman, O. Lizardo, D. Hachen, Z. Toroczka, and N. V. Chawla. Weighted reciprocity in human communication networks. *Social Networks (under review)*, 2011.
- [31] S. Wuchty and B. Uzzi. Human communication dynamics: A study of the agreement between self-reported and email derived social networks. *PLoS One, forthcoming*, 2011.