#impressme: The Language of Motivation in User Generated Content

Marc T. Tomlinson David Bracewell Wayne Krug and David Hinote
2435 N. Central Expy, Suite 1200
Richardson, TX 75080
marc@languagecomputer.com

Language Computer

Abstract. An individual's ability to produce quality work is a function of their current motivation, their control over the results of their work, and the social influences of other individuals. All of these factors can be identified in the language that individuals use to discuss their work with their peers. Previous approaches to modeling motivation have relied on social-network and time-series analysis to predict the popularity of a contribution to user-generated content site. In contrast, we show how an individual's use of language can reflect their level of motivation and can be used to predict their future performance. We compare our results to an analysis of motivation based on utility theory. We show that an understanding of the language contained in comments on user generated content sites provides significant insight into an author's level of motivation and the potential quality of their future work.

1 Introduction

Creative inspiration is only part of the puzzle to the successful completion of an endeavor. Quality work requires the setting of goals, belief in your ability to succeed, and proper feedback from the community. Community members can provide critical support at times when we are feeling down or make us to step up to a new challenge by forcing the establishment of complex far-reaching goals. Unfortunately, community members can also reduce an individual's performance through disparaging comments or lulling them into a sense of well-being. Here we present an analysis of a system for determining the motivational implicatures present within a discourse and their effect on motivating an individual to produce quality user generated content.

We cast our approach to understanding the motivational content of a communication in terms of speech acts [1], which provide a theoretical framework to explore the motivational implicatures of an utterance. In this contribution we use the term motivational act to represent utterances by individuals that either reveal their motivation for an action or affect the motivation of another individual. We classify three main types of motivational acts. The first act that we look at are comments which indicate the value of an individual's contribution, or reward (i.e. comments like "great job" indicate positive social value for the individual's work). This act is further refined into separate categories for self-directed rewards, and reward statements directed at other individuals.

The second factor looks for evidence that the individual has (or thinks they have) skill or control to act within the environment ("I know I can do this"). Lastly, we identify comments that express goals or indicators of a goal orientation in an individual.

The motivational act expressed by each comment is classified using an approach based on distant supervision [2] and twitter. We first create a language model for each motivational act based on the words contained in tweets which have been marked with a hashtag which is used by individuals on twitter to identify tweets with a motivational message.

We then show how the language models trained on twitter can be used to capture the motivational meaning of a dialogue between a contributor and commenter on a user-generated-content site. The motivational acts can be used to predict the amount of effort the individual is likely to expend on future contributions. We use the community rated quality of an individual's subsequent contribution to the site, DeviantArt.com, as a proxy for the amount of effort the individual expended on their submission. Finally, we compare the quality of the classifiers and general framework to an approach based on a utility theory view of motivation which looks at changes in the quality of submissions by that individual over time.

2 Background & Related Work

Previous research has shown how social implicatures present within a group's discussions can provide insight into the quality of a groups contributions [3]. In the spirit of dialogue acts [4–6], social acts focus on the social implicature of the statement and thus more directly relate to the social intentions and goals of individuals. Researchers have recently begun to construct and annotate social acts. Bender et al. [7] create an annotated corpus of social acts relating to *authority claims* and *alignment moves* as well as a broader selection of social acts covering managerial influence, agreement, group affordance and others. These social acts can be used to infer the quality of the work (ratings of Wikipedia articles) which a group produces based on the social interactions between group members. In contrast to that work, we focus on how comments and expressions reveal or affect an individual's motivation and their correlation with the quality of future productions.

Researchers have also examined the quality of user generated comments on blog sites using shallower stylistic features and topical information. For instance Hsu et al. [8] and Khabiri et al. [9] analyzed the popular websites Slashdot and Digg to examine prediction of the probability that a given post will receive high marks from the community. These approaches use a variety of stylistic features, such as word counts, quotes, and hyperlink counts as well as information about the reputation of the individual creating the post, temporal features of the post, and structure features pertaining to the posts location on the page. This line of work does not attempt to address the issue of predicting the quality of future contributions by the same individual.

Alternatively, a considerable amount of work has looked at predicting the popularity of user generated content on the world-wide-web. Much of the work has considered people as one-dimensional products of their social network. For example, the popularity of a given picture posted on Flickr can be derived through an examination of the social network surrounding the poster. Previous work has found that approximately 50% of all favorite markings for a given post are generated by individuals connected to the post [10]. Similarly, Szabo and Huberman[11] predict the future popularity of a YouTube or Digg contribution based on a time-series analysis of its historical popularity. Their approach requires analyzing the first few hours, in the case of Digg, or 10 days, for YouTube, and shows that after that time the content follows a fairly predictable trajectory.

In addition to work in natural language processing and social network analysis, the inspiration for our approach comes from psychological theories of motivation. The predominant theory used for understanding an individual's motivation is based on prospect theory, an extension of utility theory [12]. Kehneman and Tversky discuss three concepts that affect how an individual values a reward. 1) Reference points - rewards are valued in how far they deviate above (positive reward, gain) or below (negative reward, loss) a given reference point; 2) Loss Aversion - avoiding a loss is treated as being more important than an equivalent gain (avoiding a loss of $10 is more important than gaining $10); 3) Diminishing sensitivity - the value of a change is not linear but decreases as the point gets further from the referent ($10 to $20 is a big jump, but $1,000,000 to $1,000,020 doesn't make much of a difference). These factors can be combined to create a model of the expected utility of an action for an individual and correspondingly, an individual‘s level of motivation to achieve the reward.

Prospect theory can be related to motivation and made more concrete through theories linking goal setting as the establishment of reference points [13]. These theories suggest that an individual's motivation is a function of the expected utility of their actions. Critically, the theory suggests that goals serve to establish reference points. For example, assuming that an individual has a goal to produce great art than the value of intermediate rewards (views, favorites, positive comments by community members) will be judged using the final goal as a reference point. As they get closer to their goal chances for small positive rewards will be more motivating because they are worth comparatively more, while setbacks will be devastating. Importantly, this framework suggests that individual's that meet their goals will be less motivated on future endeavors, because they have little utility. However, in ill-defined environments, such as user-generated content sites it is difficult to know how individuals value likes from the community and what an individual is using for a reference point. In the current work we utilize a model based on prospect theory as a comparison to one based on the motivational implicatures of the language used by individuals on user-generated content forums.

3 Motivational Language Uses

Several theories exist on how different expressions of motivational factors interact with an individual's future performance. Probably the most apparent factor is the setting of goals. Goals are expressions of intentions for a change of state which could require an action on the part of the individual. Examples of goals are, "I want to finish my paper". This goal expresses an intention for a specific action but requires making inferences about the probable rewards for the individual if they are successful. In contrast, a statement such as “I want to be famous” expresses a clear expectation for a reward resulting from some series of actions, but requires inference about the details of the future actions that the individual might precipitate to achieve their goal. Does the individual want to be a rock star or a serial killer? In contrast to explicit goals stated by an individual, goals can also be inferred by other people based on an analysis of the actions carried out by the individual. For example, if an individual repeatedly demonstrates their work to a community, one can infer that the individual likely has a goal to accomplish the end-product of the activity. The setting of goals for both action and inaction has been linked to many different motivational and long-term outcomes [14, 15].

An individual’s motivation can also be signaled and inferred through their use of expressions which indicate a reward for an action. This can be derived from their communications and the communications of others. Rewards can come from an individual about their own work, “I really like my drawing”, or from other community members, “your work is top-notch”. These comments indicate a high value for the individual’s action. In contrast, comments such as “you suck” indicate a negative reward for the individual’s contribution, lowering its value and an individual’s motivation for future work. However, some care needs to be taken in this as some people actually seek out and are motivated by negative comments [16].

The last motivational act is a statement indicating control (or lack thereof) over an action. Individual‘s that feel that they have control over the outcome of an action are more motivated to perform the action [17]. Individuals express their perceptions of their control over actions through statements such as “its really easy to x” (control) or “i feel i can‘t do anything right” (lack of control).

4 Detecting Motivational Acts through Distant Supervision

Goals, rewards, and control can be expressed in a myriad of different ways in text, sometimes very clearly “I want to do better”, and sometimes only implied. This problem is analogous to that of classifying speech or dialogue acts. In this section we show how distant supervision can be used to create a language model which can identify motivational acts by a speaker.

One notorious problem with automatically recognizing the illocutionary force of an utterance is annotation of a supervised training set. Many utterances are ambiguous or signal multiple overlapping acts, this creates a notoriously difficult annotation task [18, 19]. To solve this problem we explore a novel approach for annotation through the use of a distant supervision framework [2]. Distant supervision involves the use of a small set of annotations that link to a larger knowledge base that contains noisy instances of those annotations. While the approach introduces noise into the system, the noise is mitigated by access to a very large collection of approximate annotations (in our case, millions of tweets containing these tags). Similar approaches have been used in sentiment and emotion analysis [20, 21]. In contrast, our approach learns the motivational implicature of an utterance from an individual’s use of hashtags.

Hashtags are words or phrases that are often included in tweets to signify the topic of the tweet. Some hashtags have meanings that can be derived from the words making up the tag (e.g. #mygoal – is used to express goals of an individual), while others are

Table 1. Example Hash Tags and Tweets for Goals, Control, and Rewards.

Motivational Action	Sample Tags	Sample Tweets
Goal	#goalinlife, #mywish	“3 more days of studying”
Control	#dowhatisay, #kissmyfeet	“I defy the law of gravity”
Negative Reward Self	#fml, #crap	“I just locked the keys in my car”
Negative Reward Other	#worstdriverever, #awkward	“It does make me cringe”
Positive Reward Self	#whyismile, #victoryismine	“my cats make me smile”
Positive Reward Other	#ff, #thatsbadass	“Solar panels on the white house”

related to Internet memes and require broader cultural knowledge (e.g. #fml used to express negative things happening in an individual’s own life). When a user embeds a hashtag in a tweet, twitter provides a link to a page showing a collection of all of the tweets with that hashtag. This makes hashtags very popular for researchers trying to follow trending topics on twitter. To date, most of the research has focused on linking hashtags to topics, or sentiment. However, hashtags also provide annotation for non-semantic topics, such as goals, rewards, and many other social phenomena.

We considered models for identifying goals, perception of control, and rewards based trained from hashtags that were deemed to be relevant to those phenomena. Rewards were subdivided into four categories, positive self-orientation, positive other-orientation, negative self-orientation, and negative other-orientation. The hashtags were taken from lists of trending tags as well as generated through trial and error. An initial list of hashtags was generated by a single annotator who rated each tag on a scale of 1-5 as to the overall degree that tweets using that hashtag represented the concept of interest. This list was then refined by a second annotator.

Examples of hashtags exhibiting each of the characteristics are shown in Table 1. The goals represented by these hashtags are diverse but mostly mundane, examples for #mygoal range from “3 more days of studying #iwillsurvive #4.0 #mygoal” to “Looking for a bigger house By December I wane be out this house in a bigger house #mygoal”. The English annotator identified 140 hashtags that were relevant to one of the dimensions (received a relevance score of 4 or 5 on a scale of 1-5).

We used a large collection of tweets (approximately 7.5 million). In our collection hashtags exhibiting control contained the largest number with approximately 315,000 tweets, while we only collected 110,000 tweets which were marked with a hashtag indicating positive rewards for the actions of other individuals, as shown in Table 1. For training and testing purposes we removed all URLs, hashtags, and @users from the tweets. We then discarded tweets that were less than two words long. This is very conservative, because we removed the classifier’s ability to directly learn co-occurring hashtags, however we wanted to ensure that we would minimize deficient solutions and maximize the ability of the models to transfer from twitter to other genres of text, such as web forums.

Many of the most successful approaches to dialogue act classification on text have focused on integrating multi-layer models which examine both the utterance content

Table 2. Training and testing sizes for N-gram classifiers with resultant accuracy and bias for labeling a tweet as a positive instance of the class. All test and train splits are 50/50 between positive and negative instances.

Motivational Action	# Hashtags	# Train	# Test	Accuracy
Goal	23	83,838	20,960	79.8
Control	18	153,136	38,286	70.2
Negative Reward Self	30	100,996	25,250	68.6
Negative Reward Other	47	157,582	39,396	69.6
Positive Reward Self	8	158,250	39,564	69.3
Positive Reward Other	5	103,948	25,988	78.9

and the surrounding utterances [22]. However, for this initial work we consider only the linguistic content of a single utterance.

Each motivational act detection model was trained to separate tweets that had been tagged with a hashtag that had been identified as signaling the particular motivational act from a background model that represents tweets containing hashtags which signal one of the other motivational acts. Tweets containing hashtags from multiple acts were not trained or tested on.

We utilized a language model coupled with Naive-Bayes which considered n-grams between 2-4 words in length for each of the different motivational acts. The model compared the probability of all of the n-grams from 2-4 words long in the tweet given the motivational act to the probability of the n-grams given a background model. For evaluation purposes we considered the act to be present if the sequence of words was more likely given the act than in the background model, but for downstream usage, the model output the odds of the act compared to the background model.

The accuracy of the resultant classifiers (show in table 2 suggest that they are adequately capturing the differences between the categories, though there is some obfuscation of the true validity of the labels due to noisy use of the tags by individuals. Inspections of tweets with the labels suggest that many times the labels are used sarcastically. Anecdotally, we also examined a list of the top hashtags associated with instances labeled by our approach and found good generalization to novel hashtags. We looked at a list of the hashtags based on the average confidence of the labels being applied to the tweets containing those tags, we found many reasonable candidate tags. For example, tweets containing the hashtags #day1 and #day2 were among the most likely to be labeled as exhibiting a goal. This suggests that a model which spiders out proposing new hashtags based on co-occurrence and is able to then incorporate those instances would work very effectively and allow for fine tuning of the model.

The classifiers created above are capable of identifying the motivational signals found within an individual's language. In the next section we show how those classifiers can then be transferred to a novel domain and allow us to understand an individual's motivation by examining a complete discourse.

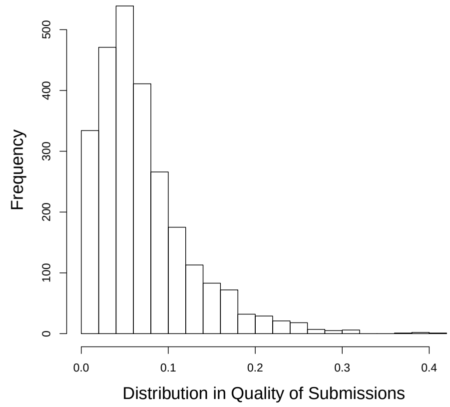

Fig. 1. Distribution of quality on DeviantArt. Quality is measured as the probability of an individual favoriting a contribution by an artist.

5 Modeling Motivation

In the previous section we showed how an individual utterance can be understood in terms of its motivational meaning. In this section we show one way to build a model of the amount of effort an individual is willing to expend on a future endeavor by examining the motivational acts that the individual and other conversational participants use. For this effort we look at modeling the quality of an individual's contribution to the user generated content website DeviantArt.com.

5.1 Data

We collected and analyzed date from the user generated content site DeviantArt.com. DeviantArt is a community web-forum on which individuals post art-work that they created and receive comments on their posts. The art is very diverse including pictures, cartoons, and needlework, though as the name suggests many of the works are avant-garde. The website allows for artists to post their work and a forum for discussion between the artist and the community members about each contribution. Community members can identify their favorite pieces. The site is very well established and boasts over 2 million subscribers.

The data contained on DeviantArt and related sites is very valuable, because it provides easily inferable goals (respect from community members) as well as explicit statements that can improve these inferences, clear actions (production of art) by community members, and clear rewards for actions (community response, favorites, views). Most importantly, it has an easily inferred timeline of related actions which receive clear feedback at most states.

DeviantArt provides several different ways to assess the success of each piece of art. The website keeps track of the number of views that each piece of art receives, the number of comments it receives, and the number of times that individual's tag it as a favorite. A large percentage of the work examining user generated content looks at predicting the number of views that an item will receive. While previous work has reported that views are not well correlated with comments or favorites on Flickr [10], the views and favorites on DeviantArt are highly correlated ($\rho = .84, p < .001$, because of the distribution of these measures we are reporting the non-parametric rank-order correlation). Also in contrast to Flickr, comments and views, and comments and favorites are only weakly correlated ($\rho = .24,\ \rho = .22, p < .001$, respectively). This could be due to differences between the two communities, or differences in how browsing of new pieces is supported by the website.

Most previous work has looked at predicting the popularity of a piece, in contrast we break the trend by looking at a hybrid measure of the probability that an individual favorites a piece given that they viewed it. We feel that this measure provides a better characterization of the overall quality of a contribution. This is supported by an analysis of the correlation between the popularity of a piece and our measure. The probability that a piece is favorited given that it was viewed is the least correlated measure with the number of views; the number of comments on a piece is second-least correlated. For this paper we refer to our aggregated measure as the quality of a submission, but more research should be done to better characterize the differences in these measures.

To generate our data we sampled from the artists posting content on DeviantArt generating an initial list of 1103 artists. Each artist contributed, on average, 20 contributions for a total of 21,420 pieces of art. The mean number of views per contribution was 1545, while the median was 248. A preliminary analysis of the data showed that the variance in the quality of an artist's contributions was much higher for new artists than for experienced artists. The data suggest very interesting variations among individuals of different popularity and experience levels, but based on this preliminary analysis we restricted our analysis to those contributions where the artist had at least 1000 views and was at least the 20th post by the author. By restricting the data to those with 1000 views we also reduce the noise in the quality measure because of the size of denominator. This reduced our data set to 101 artists, with 2,059 contributions, and 127,622 comments on the artwork.

5.2 Using Comments to Infer Motivation

To understand how comments on a web forum can reflect an individual's level of motivation we created a model, the goals, reward, and control model (GRC), which determined the extent to which an individual's conversation with group members exhibited qualities of a motivated individual based on the motivational acts used within a conversation. We compared this model to a base-line model which used sentiment terms (Sentiment Model) instead of motivational acts and also to a second model which looked at predicting the motivation level of an individual based on the change in quality over time (Utility Model). Finally, we show how the GRC model can generate improved predictions through use of the utility model.

Table 3. Example Goals, Control, and Rewards found on Flickr.

Motivational Action	Sample Comments
Goal	"Hoping to finish this in time for Easter."
Control	"Great details"
Negative Reward Self	"but it is very poor quality"
Negative Reward Other	"It looks a little bit flat with not much contrasts "
Positive Reward Self	"Thank you so much!"
Positive Reward Other	"Stunning photo! Congrats"

The GRC model uses an analysis of the artist's comments and the comments by other conversational participants to construct a model of the artist's motivation. Each of the comments were first identified as being made by the artist or by another individual. They were then labeled using the language models discussed in section 4 for each of the six motivational acts:Goal, Control, Negative Reward Self, Negative Reward Other, Positive Reward Self, Positive Reward other. Examples are shown in Table 3 For each comment page the number times that language was used which exhibited each act were aggregated separately for the artist and community members. These twelve features were weighted and combined linearly, $\sum_{1..12} \beta_i F_i$.

Our base-line, sentiment model utilized an equivalent linear regression procedure except in place of the motivational acts we examined the presence or absence of positive and negative sentiment terms in each communication by the artist and the community. We used SentiWordNet [23] taking all terms with a positivity or negativity higher than .5.

Likewise, the utility model uses a linear combination of two factors to predict the quality of a future contribution. Our first factor is the expected utility of the next contribution, we set this as an exponential function of the quality of the current submission, Q_t^λ. We used a λ value of .3 to encode the diminishing sensitivity to higher rewards, though testing showed less than a 1 percent difference for realistic values of λ and a similar patterning of results. The second factor in the model accounts for goal related behavior (establishing the current piece as a goal) and adjusts the utility based on $Q_t - Q_{t-1}$ or $\Delta Q_{t-1,t}$. Individuals that accomplish their goals, $Q_t > Q_{t-1}$ are less motivated. For the results we split this model into one utilizing the first factor (quality) only and one showing the first and second factor combined.

The last model combines the GRC and the Utility Model to look at the interaction between the expected motivation of the individual based on their previous rewards and the way in which the individual and group is discussing their rewards, goals, and perceptions of control. This model replaces $\Delta Q_{t-1,t}$ from the utility model with the linguistic features, F_i.

5.3 Modeling Motivation and Results

For each of the four models we looked at two separate predictions. The first is how well we can predict the quality of the next submission by an artist, Q_{t+1}. Thus for the GRC

Model we find the maximum likelihood estimate minimizing the error for the following equation

$$Q_{t+1} = \beta_0 + \sum_{1..12} \beta_i F_i.$$

using least squares and correlate our predicted value $\hat{Q}_{t+1}$ with the actual value Q_{t+1} on held out test data.

We utilized a 100-fold cross-validation procedure to find the correlation between the predicted quality of the next submission and the actual quality for each artist. The data for each artist was distributed randomly across the folds. The results are shown in the first column of Table 4. The GRC model achieves a correlation of .33, which is significantly greater than the baseline sentiment model. The .51 obtained by the utility model shows the correlation between the quality of a previous post and the next post, by adding in the difference in quality between the post at time t and $t-1$ we see another boost in performance to .53. Finally, the combined model achieves a correlation of .52, which is slightly better than the Quality based Utility model.

The above results suggest that the language individuals use reflects the quality of their future contributions. However, it is quite clear that the quality of the current submission is the best predictor of the quality of an individual's next submission. It could be that the comments in the post are only associated with the current quality and not the individuals future motivation, we want to measure changes in an individual's motivation over time.

One approach to measuring the change in motivation across time would be to consider if the quality of an individual's next submission is higher or lower than the quality of their current submission. However, this measure is subject to regression to the mean. Assuming that the quality of an individual's submission is normally distributed, if an individual produces an above average or below average submission, the next contribution is more likely to be below or above that contribution as a simple by-product of probability theory and not their level of motivation. Instead, we can look at the difference between the quality of a contribution at t+1 and its expected value based on the mean level of quality for the individual.

Our second prediction tests categorically, utilizing logistic regression, whether an individual's next post will be above or below average, defined as the moving average of the quality of their last 20 posts $\bar{Q}_t$. The results shown in the second column of Table 4 suggest that the model based on motivational acts performs at a level more similar to that based on utility theory. Additionally, the combined model performs better than either individual model. Critically, this second analysis shows that an individual's language usage reflects their level of motivation for subsequent contributions, successfully predicting whether or not an individual will contribute content of unusually high or low quality compared to their average level of performance.

6 Conclusion

This paper serves as an initial demonstration of an approach for identifying an individual's motivation for future work. Our analysis was conducted by examining the motivational meaning of the language that the individual used in comments about their work

Table 4. Results for predicting the future quality of a contribution based on analysis of the users previous contributions and comments on those contributions. Differences of more than .01 are significant at $p < .05$ according to a paired t-test across testing folds.

Method	Q_{t+1} Correlation	$Q_{t+1} > \overline{Q}_t$ Classification Accuracy
Majority Class	NA	.54
Sentiment	.08	.55
GRC	.33	.58
Utility (Quality Only)	.51	.61
Utility & Difference	.53	.63
GRC + Utility	.52	.65

and the language used by community members about the work. In addition, we showed that distant supervision of natural language classifiers can be used to identify, not just sentiment or semantics, but the language's effect on an individual's psychological state. For example, the language models can reveal their goals, perception of control over a situation, and the rewards that motivate them.

It is important to remember that our results show a link between language usage and future performance, but the results do not imply causality in a particular direction. The categorical prediction identifies individuals that will perform above average on their next submission, but the language could be a by-product of their motivation instead of the cause of their motivation. It will be up to future work to look at this in more detail and attempt to separate linguistic expressions that inform motivational states and those that change motivational states. In particular, to look at how the pattern of interactions between posters and community members can reveal the motivational level of the individual. Our focus in this contribution was in demonstrating a system for automatically detecting language uses that provides insight into an individual's level of motivation and that reflect the amount of effort an individual is likely to apply to their future performances.

Acknowledgment

This research was funded by the Intelligence Advanced Research Projects Activity (IARPA) through the Department of Defense US Army Research Laboratory (DoD / ARL). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoD/ARL, or the U.S. Government.

References

1. Searle, J.R.: Speech Acts: An Essay in the Philosophy of Language. Volume 0. Cambridge University Press (1969)

2. Mintz, M., Bills, S., Snow, R., Jurafsky, D.: Distant supervision for relation extraction without labeled data. In: Proceedings of ACL-IJCNLP. Number 2005 (2009)
3. Bracewell, D.B., Tomlinson, M.: The Social Actions of Successful Groups. In: IEEE Sixth International Conference on Semantic Computing. (2012)
4. Core, M.G., Allen, J.F.: Coding Dialogs with the DAMSL Annotation Scheme. In Traum, D., ed.: AAAI Fall Symposium on Communicative Action in Humans and Machines, AAAI, American Association for Artificial Intelligence (1997) 28–35
5. Stolcke, A., Shriberg, E., Bates, R., Coccaro, N., Jurafsky, D., Martin, R., Meteer, M., Ries, K., Taylor, P., Ess-Dykema, C.V.: Dialog Act Modeling for Conversational Speech. In: Applying Machine Learning to Discourse Processing, AAAI Press (1998) 98–105
6. Bunt, H.: The semantics of dialogue acts. International Conference on Computational Semantics (2011) 1–13
7. Bender, E., Morgan, J., Oxley, M., Zachry, M., Hutchinson, B., Marin, A., Zhang, B., Ostendorf, M.: Annotating social acts: Authority claims and alignment moves in wikipedia talk pages. ACL HLT 2011 (June) (2011) 48
8. Hsu, C.F., Khabiri, E., Caverlee, J.: Ranking Comments on the Social Web. 2009 International Conference on Computational Science and Engineering (2009) 90–97
9. Khabiri, E., Hsu, C., Caverlee, J.: Analyzing and predicting community preference of socially generated metadata: A case study on comments in the digg community. AAAI Conference on Weblogs and Social Media (2009)
10. Cha, M., Mislove, A., Gummadi, K.P.: A measurement-driven analysis of information propagation in the flickr social network. Proceedings of the 18th international conference on World wide web - WWW '09 (2009) 721
11. Szabo, G., Huberman, B.A.: Predicting the popularity of online content. Communications of the ACM **53**(8) (2010) 80–88
12. Kahneman, D., Tversky, A.: Kahneman & Tversky (1979) - Prospect Theory - An Analysis Of Decision Under Risk.pdf. Econometrica **47**(2) (1979) 263–291
13. Heath, C., Larrick, R.P., Wu, G.: Goals as reference points. Cognitive psychology **38**(1) (February 1999) 79–109
14. Albarracin, D., Hepler, J., Tannenbaum, M.: General Action and Inaction Goals: Their Behavioral, Cognitive, and Affective Origins and Influences. Current Directions in Psychological Science **20**(2) (April 2011) 119–123
15. Locke, E.A.: Toward a theory of task motivation and incentives. Organizational behavior and human performance **3**(2) (1968)
16. Finkelstein, S.R., Fishbach, A.: Tell Me What I Did Wrong: Experts Seek and Respond to Negative Feedback. Journal of Consumer Research **39**(1) (June 2012) 22–38
17. Ajzen, I.: The Theory of Planned Behavior. Organizational Behavior and Human Decision Processes **50** (1991) 179–211
18. Fang, A., Bunt, H., Cao, J.: Collaborative Annotation of Dialogue Acts: Application of a New ISO Standard to the Switchboard Corpus. EACL 2012 (2012) 61–68
19. Bracewell, D.B., Tomlinson, M.T., Brunson, M., Plymale, J., Bracewell, J., Boerger, D.: Annotation of Adversarial and Collegial Social Actions in Discourse. In: 6th Linguistic Annotation Workshop. Number July (2012) 184–192
20. Davidov, D., Tsur, O., Rappaport, A.: Enhanced Sentiment Learning Using Twitter Hashtags and Smileys. In: Coling. Number August (2010) 241–249
21. Roberts, K., Roach, M., Johnson, J., Gurthrie, J., Harabagiu, S.M.: Empatweet: Annotating and detecting emotions on Twitter. Proceedings of LREC2012 (2012) 3806–3813
22. Petukhova, V., Bunt, H.: Incremental dialogue act understanding. IWCS '11 Proceedings of the Ninth International Conference on Computational Semantics (2011) 235–244
23. Baccianella, S., Esuli, A., Sebastiani, F.: SentiWordNet 3.0 : An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining. In: LREC. Volume 0. (2010) 2200–2204