Capturing Cultural Differences in Expressions of Intentions

Marc T. Tomlinson   **David B. Bracewell**   **Wayne Krug**
Language Computer
Richardson TX 75080
`marc, david, wayne@languagecomputer.com`

Abstract

The intersection of psychology and computational linguistics is capable of providing novel automated insight into the language of everyday cognition through analysis of micro-blogs. While Twitter is often seen as banal or focused only on the *who*, *what*, *when* or *where* tweets can actually serve as a source for learning about the language people use to express complex cogntive states and their cultural identity. In this contribution we introduce a novel model which captures latent cultural dimensions through an individual's expressions of intentionality. We then show how these latent cultures can be used to create a culturally-sensitive model which provides enahnced detection of signals of intentionality in tweets. Finally, we demonstrate how these models reveal interesting cross-cultural differences in the goals and motivations of individuals from different cultures.

1   Introduction

Social media platforms have enabled new forms of discourse and have also provided enormous quantities of data on these communications. For instance, the popular microblogging service Twitter provides an exceptionally useful source of user-generated content which has attracted considerable interest from researchers in computational linguistics (Ritter et al., 2009; Gimpel et al., 2011). Most of the language processing on tweets has involved the identification of sentiment (Davidov et al., 2010), summarization (Sharifi et al., 2010), conversational models of Dialogue acts (Ritter et al., 2009), or lexical and semantic processing. In this effort we expand on these previous approaches and show how individuals express their cultural identity through expressions revealing their intentionality towards events and provide a way of capturing this information.

We define intentionality as the amount of effort an individual is willing to expend to achieve a goal(Ajzen, 1991). Goals represent future states or events which an individual wishes to happen. Accordingly, intentions are goals for which an individuals is willing to expend at least some minimal amount of effort to bring about. While people express goals throughout the day, intentions are the goals that they are willing to follow through with. Identifying when a goal is actually an intention requires the successful recognition of many distinct cognitive factors that can be revealed through the individual's use of language.

There is a long history of studies that have worked towards identifying a set of factors that underly an individual's intentions (Ajzen and Fishbein, 1977; Ajzen, 1991; Malle and Knobe, 1997; Sloman et al., 2012) of which, the setting of goals is one important factor. These studies have concentrated on identifying the factors that affect an individual's motivation. The studies have also identified a set of factors that people use to gauge the intentionality of other individuals. However, these factors have always been manually identified by an expert from an individual's speech or writing. It is not clear that these features can actually be detected automatically in language.

Intentions have also been considered in computational linguistics. In their seminal work entitled, "Attention, Intentions, and the Structure of Discourse" (Grosz and Sidner, 1986), Grosz and Sidner point out the fundamental role of intentions and their effect on the theory and processing of discourse structure. They even define a set of intentions that can be held by individuals that are relevant to discourse theory. In contrast, we focus on understanding intentions outside of the discourse. In addition, we work with a more general definition of intentions taken from psychology, defining intentionality as the amount of effort an individual is willing to expend to achieve a goal.

Culture refers to the set of beliefs, norms, and customs shared by a group of people. Beliefs and culture are inseparably tied to intentions and language (Ajzen and Fishbein, 1977; Tomasello et al., 2005). Culture affects an author's proclivity to have a particular intention, for example Hofstede's dimension of power distance (Hofstede, 1980) would suggest that individuals from high power-distance cultures have a lower likelihood of performing

actions with the intention of overriding the actions of an individual of higher status. Culture can also affect the way in which individuals reason about other agents' intentions and the set of actions that are used to realize an individual's intentions. While considerable work has looked at the link between cultures and intentions, here we show how a latent representation of an individual's culture derived from their intentions can be utilized to explore the intersection between culture and intentions using the vast amount of written expressions present on Twitter.

In this contribution, instead of focusing on the discourse meaning of intentions, we look at how personal intentions can be understood through Twitter posts by focusing on the language of those posts contain. We briefly discuss previous work showing how it is possible to capture language that reveal cognitive factors of intentionality which could be used to capture broader intentions. Critically, we then augment the models of the cognitive factors of intentionality by accounting for the culture of the authors on Twitter. Twitter contains an immense number of authors covering a variety of different cultures definable at different levels, for example women, college-students, or fitness buffs.

We have evaluated the models on a very large set of over 7.5 million tweets which cover a sampling of Twitter from early 2011 to the middle of 2013. Our sample includes just over 900,000 authors. We found very promising results for identifying the factors of intentionality, but by considering culture we were able to provide a significant improvement of those results. We have shown that cognitive factors of intentionality, including goals, control and skill, and rewards can be recognized through the use of simple language models. Similarly, our cultural models were based on traditional techniques for latent variable modeling through principal component analysis enabling an understanding of the cultural distribution of intentions.

The remainder of the paper is organized as follows. We first present the cognitive factors of intentionality that we have used for this contribution. We then present a new cultural model of authors on Twitter and compare it to existing approaches in the literature. We then present a series of models which capture the cultural variation of the cognitive factors of intentionality. Finally, we present a look at some of the cultural differences identified through our approach.

2 Factors of Intentionality

While there are numerous factors that affect an individual's intentionality (Ajzen, 1991; Malle and Knobe, 1997; Sloman et al., 2012), in this contribution we focus on investigating the most historically central factors: goals, perceptions of control, and rewards. Below we provide brief examples of the three factors before detailing our approach for identification of latent cultures.

2.1 Factor 1: Goals

The first factor that we consider is evidence that an individual has a goal. Goals are expressions of a desire for a change of state or rewards which could require an action on the part of the individual. The setting of goals for both action and inaction have been linked to many different motivational and long-term outcomes (Albarracin et al., 2011; Locke, 1968). Examples of goals are

(1) I want to finish my paper

(2) I want to be famous

The first example of a goal expresses an intention to perform an action which could result in a positive reward for the individual, however it doees not mention the reward. In contrast, the second example expresses a clear expectation for a reward (fame), but does not describe the actions that will lead to that reward. Does the individual want to be President or the next Kardashian? Additionally, in contrast to explicit goals stated by an individual, goals can also be inferred by other people based on an analysis of actions (perceived intended events) carried out by the individual. For example, it is presumed that an individual has a goal to win the lottery when they buy a lottery ticket, or that the occupants of a car full of beach toys is headed to or from a beach. Goals represent the factor that has seen the most recent attention in terms of the creation of automatic methods for their recognition (Chen et al., 2013; Banerjee et al., 2012).

2.2 Factor 2: Perception of Control

Intentions are revealed not just through goals, but also through words expressing skill or a level of control. Individuals that feel that they have more control over a situation will expend more effort on their actions (Ajzen, 1991). Individuals are also perceived by others as having greater intentionality for actions that they have control over or exhibit skill at. We considered multiple ways in which an individual can express their perceived control over an event, subdividing this factor into three sub-factors. The first sub-factor captures expressions which indicate skill.

(3) Just helped some guy push his gas-less car to the garage #iamwoman #hearmeroar

Table 1: Example Hash Tags and Tweets.

Cognitive Factor	Sample Tags	Sample Tweets
F1: Goal	#goalinlife, #mywish	"3 more days of studying"
F2: Control	#dowhatisay, #kissmyfeet	"I defy the law of gravity"
F2: Skill	#madskillz, #iamapro	"you are flat out amazing to watch"
F2: Lack of Control	#oops, #cantstop	"cannot believe I said that"
F3: Negative Reward Self	#fml, #crap	"I just locked the keys in my car"
F3: Negative Reward Other	#worstdriverever, #awkward	"It does make me cringe"
F3: Positive Reward Self	#whyismile, #victoryismine	"my cats make me smile"
F3: Positive Reward Other	#ff, #thatsbadass	"Solar panels on the white house"

The second sub-factor captures expressions of control.

(4) I'm in control here!

The third sub-factor captures expressions of lack-of-control.

(5) i'm a little nervous for tomorrow

While several linguistic theories exist that could be utilized to create systems detecting control, such as agency (Dowty, 1991), there is no prominent work on automatically identifying control directly in an individual's expressions.

2.3 Factor 3: Reception of Rewards

Intentions can also be inferred when an individual receives a reward.

(6) I'm so proud of what I did

(7) Your work sucks!

Rewards can be positive (increasing the likelihood of the action being repeated, Example 6) or negative (decreasing the likelihood of the action in the future, Example 7). In addition, rewards can come from the individual (self-directed rewards, Example 6) or from other individuals (other-directed rewards, Example 7). This establishes four sub-factors for rewards. Knowing that an individual received a reward increases the likelihood that they had effortful participation in the event. In addition, evidence of negative rewards are strongly inferential for intentionality (Knobe, 2003). Interpretations of rewards are very culturally sensitive. For example, a comment such as "That is disgusting" would have a good chance of being interpreted as a positive reward when it was made as a comment to a user-generated contribution on the website DeviantArt.com. Additionally, the effect of rewards on motivation is not always clear-cut. Experts seek out and are actually motivated by criticism (Finkelstein and Fishbach, 2012).

2.4 Linking Hashtags and Factors of Intentionality

The factors and sub-factors described above capture expressions which can be used to infer an individual's intentionality towards a future action. In Tomlinson et al. (2014) we showed that it is possible to link particular hashtags used by people on Twitter to these cognitive factors. Our approach utilized two annotators. The first annotator, through trial and error, identified a large number of potential candidate tags for each sub-factor. The annotator then rated each hashtag for how well tweets containing that hashtag exhibited each sub-factor (on a scale of 1-5). The second annotator then separately rated each tag which scored a 4 or 5. The two annotators had an agreement rate of 87%. 178 tags in all were agreed to be a 4 or 5 by both annotators and considered representative of the particular sub-factor. Examples of the hashtags utilized and tweets with those tags are shown in Table 1. The tweets have been modified slightly to preserve anonymity.

3 Identification of Latent Cultures

In the preceding section we discussed examples of goals, control, and rewards, and discussed how hashtags are used on Twitter to mark a tweet expressing one of these factors. Some of these examples require cultural knowledge in order to correctly interpret. In this section we present the latent model of culture that is used for learning the cultural specific expressions of the factors of intentionality.

3.1 SVD-Model of Culture

A considerable amount of work has demonstrated how particular social characteristics of individuals can be identified on Twitter, such as gender, age, and political orientation (Zamal et al., 2012; Pennacchiotti and Popescu, 2011). While superb results can be obtained for identifying these characteristics of authors using a complex set of features, this approach does not neccesarily allow for generalization to other data sets. Therefore we settled on an approach utilizing a specially trained latent variable model. Instead of utilizing Latent-Dirichlet Allocation (LDA, Blei, Ng, and Jordan, 2003) as Pennacchiotti and Popescu we utilized a spectral analysis based on singular-value decomposition (SVD). This approach has been shown to be generally superior to LDA on the domain of topic modeling (Chen et al., 2011), but has not been tested for cultural modeling.

3.1.1 Data

We randomly sampled 1.6 million tweets from a Twitter dataset that had been generated by retrieving tweets that carried at least one of the hashtags linked to a cognitive factor of intentionality (and other posts by that author). In addition, we restricted the set to authors for which we had at least 20 posts in our dataset. For this dataset, all of the markup was left in the tweet (e.g. hashtags, urls, etc.).

3.1.2 Model

From our dataset we created a set of documents, $D = \{a_1, a_2, \ldots, a_A\}$. Where each a_i represents the entire collection of tweets for a single author that contain mentions of goals, skill/control, or rewards. This set of documents contains N words and hashtags. We then create a matrix, $X \in \Re^{N,A}$, where each author represents a row in the matrix and the columns are the number of times that the corresponding word or hashtag was used by that author. Then we perform a singular value decomposition of the matrix to solve

$$X = VSC^T \tag{1}$$

Where S is a $k \ x \ k$ matrix whose off-diagonal entries equal 0 and the on-diagonal entries are the k singular values for the matrix X. For our approach we set k equal to 100. V represents a mapping of the words into our reduced space $\Re^{n,k}$, and $C \ \Re^{i,k}$ contains a weighting for each author with respect to the k^{th} latent cultural dimension. The cultural model can be used to identify the culture of an unseen author through the creation of a projection matrix, P.

$$P = VS^{-1} \tag{2}$$

This matrix projects the tweets that make up the author into our latent cultural space C. This allows us to map each author in our complete data set into our latent space which can then be used for training and testing. The latent cultural space can be used to characterize the culture of an author as a distribution over the dimensions. Below we evaluate our latent cultures on the shared dataset provided by Zamal et al. 2012.

3.2 Evaluating the Latent Cultures

Culture is a system of shared beliefs and actions. Culture is often shared between individuals based on social similarity, this can be within a language, nation, gender, age-group or other social distinction. Thus, being able to identify an individual's culture should facilitate detection of socio-demographic information. To test this we looked at using the latent cultural dimensions to predict socio-demographics on Twitter. We looked at the systems ability to identify gender (male vs. female), age (young vs. old), and political orientation (Democrat vs. Republican) of individuals based on their exhibition of particular latent cultural dimensions. In this model we first represented an individual's tweets as a distribution over the latent dimensions. We then utilized two different statistical approaches to find associations between particular dimensions and the relevant socio-demographic information. For a comparison, we tested our SVD-culture model against a similarly trained LDA model and a model based on n-grams.

3.2.1 Data

We utilized the publicly available dataset from Zamal, Liu, and Ruths (2012). The dataset consisted of Twitter user names and associated meta-data identifying their gender (Male or Female), age (two classes, young and old), and political orientation (Liberal or Conservative). Unfortunately, many of the identified tweets were no longer available from the Twitter API, but we successfully retrieved 2.6 million tweets from authors identified in the dataset with 310 users identified for gender, 320 identified for their age, and 380 for their political affiliation. The tweets in our dataset are substantially different from the original dataset because of the time over which they were collected. Zamal, et al.'s tweets were from 2012 and before, whereas our tweets covered much of 2013. This suggests that comparisons of the raw numbers should be made with caution, particularly in the political area.

Table 2: Results for identifying user demographics based on latent cultural dimensions compared to linguistic style and an ensemble method utilized by Zamal et al (2012).

		Zamal et al.	N-Grams	LDA	SVD
	N	F	F	F	F
Gender	310	.80	.57	.71	.70
Age	320	.75	.63	.66	.67
Political Orientation	380	.89	.73	.66	.68

3.2.2 Modeling & Results

To provide a comprehensive view of the strengths and weakness of our approach we compared several models for their ability to correctly predict the cultural demographics of individuals on Twitter. We first established a base-line model which was an n-gram language model created from the language used by each individual in their tweets. This model learned to identify the cultural demographics based on the frequency with which individual's in that demographic used sequences of words, called n-grams. This approach is consistently ranked as one of the single best approaches to authorship identification and performs well on a large variety of datasets.

We also tested the SVD-Culture model introduced above on this dataset. For this experiment, we trained a logistic-regression based classifier to identify the demographic information of an author based on the vector created by projecting that author into our latent space.

Finally, to look for a difference in the performance between an SVD-based latent representation and one based on LDA (Pennacchiotti and Popescu, 2011), we also trained and tested an LDA-based Culture model. The model was trained on the same data as the SVD-based model and utilized the same number of dimensions.

All of the models were tested and trained utilizing 10-fold cross validation. It is very important to point out that the data sets used to generate the underlying latent representational models did not include any of the tweets from the data used for the 10-fold cross validation. That data was only utilized for the supervision of the logistic regression.

The results of the base model, the SVD model, the LDA model, and the original results presented by Zamal et al. (2012) are shown in Table 2. The latent models are clearly superior to the language model, on average outperforming it by a significant margin of 4%. As expected, the SVD-based model does outperform the LDA model on average, though it is only by 1%, on average.

The strength of this approach is in its simplicity. The latent cultural dimensions have been learned on a wholly different dataset than that used for testing, this supports good generalization performance. While the latent SVD-cultural model does not reach the performance of the system created by Zamal et al. (2012). Zamal et al.'s results were obtained using a plethora of different feature types, which were specifically trained to solve each individual problem. As pointed out in Cohen and Ruths (2013) this causes some issues on transfer to a novel dataset, because the selected features were not representative of differences between liberals and conservatives in the second dataset. In contrast, we suggest that the latent cultural model learns a more general representation utilizing only the set of features provided by the underlying latent cultural models, which were not trained on any of the data in the test set. Additionally, the latent SVD model is easy to implement and train.

Importantly, these results indicate that the latent cultural dimensions capture similarities in the ways in which individuals of similar socio-demographics express themselves on Twitter. The model is able to easily identify the gender, age, and political affiliation of individuals based on their tweets. In the next section we show how we can utilize these latent cultural dimensions to facilitate learning of expressions conveying factors of intentionality.

4 Cultural Sensitive Identification of Cognitive Factors of Intentionality in Language

Recognizing language that expresses factors of intentionality is complicated because of the wide variety of ways in which they can be expressed as shown in the examples in the previous section. While some work has explored automatic goal recognition, most recently by (Chen et al., 2013) and (Banerjee et al., 2012), little work has been done automatically characterizing the other factors, though work in detecting social implicatures in language is similar (Bracewell et al., 2012b). We first present a general framework for learning to model the content of tweets that express a given factor from our cognitive model, we then show how this approach can be enhanced with the addition of latent cultural dimensions.

4.1 Culture Agnostic Model

Here we introduce the *General Model* that serves as the basis for the culture specific models. It is so named because it applies to all cultures. We utilized an n-gram based language model to identify the factors in tweets.

We first constructed a vocabulary of all n-grams between 2 and 4 words in length. Each tweet, j, which is labeled with a hashtag linked to a sub-factor f, is represented as a vector, X^j. Entries in X^j correspond to the number of occurrences in the tweet of the i^{th} n-gram from the vocabulary. We examined two different mathematical approaches to modeling the cognitive factors to gain a better understanding of the problem.

The first approach utilized a Naive-Bayes based classifier (NB) where

$$p(F = f|X) = \frac{p(F = f) \cdot p(X|F = f)}{p(X)} \tag{3}$$

The second approach utilized an L2-loss logistic regression model (L2):

$$p(F = f|X, W) = \frac{1}{1 + exp(w_0 + \sum_i w_i X_i)} \tag{4}$$

In which the weights, W, are learned by maximizing Equation (4)

$$\sum_j^{m_f} log \, p(y^j|X^j; W) - \alpha ||W||_2^2 \tag{5}$$

where m_f represents a balanced training set created by randomly sampling the training tweets that are tied to sub-factor f and an equal number of tweets that express one of the other factors. For solving the maximization problem we utilized the LibLinear package (Fan et al., 2008).

4.2 Culture Sensitive Models

We compared two different methods for integrating the culture information from the SVD-based culture model into the models for identifying the cognitive factors of intentionality. Both models assume that the authors have been partitioned into a set of cultures, L, but differ in their modeling of the link between language and cognitive factors.

In order to identify the cultures of the authors we utilize a clustering of the latent dimensions produced by the SVD model, a spectral clustering (Kannan et al., 2004). We utilized a simple hierarchical clustering that capitalizes on the y largest singular values. We create a set of hierarchical clusters based on a median split of each of the first y columns in our latent space. When $y = 1$ we have two clusters where the authors have been split based on the median value of the first latent dimension, with $y = 2$ each cluster is then independently split by the author's value along the second latent dimension, giving four clusters, and so on.

4.2.1 Culture-Specific Model

Our first method, which we call the culture specific model uses a separate model of each factor for each latent culture, $l \in L$. We first identify a tweet x, as belonging to a given culture, l. We then determine whether or not the language it contains expresses a particular cognitive factor based on

$$p(F = f|x_l, L_l) \tag{6}$$

To learn the function we utilize a linear classifier, Logistic-Regression with an L2 regularization term, and limit the training data to authors that belong to the particular culture.

4.2.2 Joint-Culture Model

Our second model, which we call the joint culture model utilizes an ensemble based approach. For each tweet, x_l, we calculate both a culture specific view of the language in the tween $p(F = f|x_l, L_l)$ and a culture agnostic view $p(F = f|x_l)$, taking the classification is that is most confident. This joint approach utilizes the culture-agnostic model to smooth deficiencies caused by insufficient culture-specific data.

4.2.3 Number of Cultures

We explored settings of $y = \{2, 3, 4, 5\}$ latent dimensions which equates to $\{2, 4, 8, 16, 32\}$ latent cultures. Authors are first split according into their cultural group and then tweets from each culture are broken into a training and testing set. Because of the amount of data we utilized only a 5-fold cross validation procedure. In addition, we also tested a random culture model that randomly assigned authors to cultures instead of utilizing the spectral clustering. When creating these random cultures we balanced the number of authors in each random culture with the corresponding spectral cultures.

Table 3: Accuracies for modeling each sub-factor of intentionality. L2 represents results obtained using an L2-regularized linear regression, NB represents naive-Bayes, #Cultures signifies the number of latent cultural dimensions used for clustering.

| #Cultures | General | | L2 - Culture Specific | | | | L2 -Joint Culture | | | | |
	NB-0	L2-0	2	3	4	5	2	3	4	5	NB-5
F1:Goals	79.8	80.9	81.1	80.9	79.8	78.8	82.1	**82.2**	82.1	82.0	79.1
F2:Control	70.1	75.5	75.5	75.3	74.4	73.8	76.4	**76.4**	76.4	76.4	72.3
F2:Lack of Control	69.1	73.7	75.2	74.9	74.1	72.9	**75.9**	75.7	75.8	75.6	71.6
F2:Skill	73.2	76.2	77.6	77.0	76.4	75.5	78.2	**78.2**	78.1	78.1	75.3
F3:Positive Other	78.3	82.9	84.3	84.1	83.7	83.4	**84.5**	84.4	84.4	84.5	81.9
F3:Positive Self	66.0	69.1	70.6	70.3	69.4	68.7	71.3	71.3	71.3	**71.4**	68.4
F3:Negative Other	68.7	72.3	73.6	73.4	72.5	71.6	74.1	74.0	**74.1**	74.0	70.8
F3:Negative Self	69.3	72.4	73.6	73.3	71.9	71.3	**74.4**	74.3	73.9	73.7	71.1

4.3 Data

Testing was done on a large number of tweets (7.5 million) that contained tweets from individuals that used any of the representative hashtags. In our collection hashtags exhibiting the sub-factor of control contained the largest number with approximately 575,000 tweets, while we only collected 110,000 tweets which were marked with a hashtags indicating positive rewards for the actions of other individuals. 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 approach is 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.

4.4 Results and Discussion

The accuracy of the classifiers for identifying each sub-factor are shown in Table 3. The accuracies reflect the classifiers ability to separate tweets that have a hashtag representing the given sub-factor from those that do not. The results suggest that all of the models are adequately capturing the differences between the cognitive factors. On average, the logistic regression based classifier achieves a 3.5 percent advantage in accuracy over the Naive-Bayes model, showing a clear advantage for the improved feature selection of the L2-loss logistic regression. Both models required a similar amount of time to train and test.

To conserve space Table 3 shows only the results for the 5-dimension Joint Culture Naive-Bayes model. The results for the Naive-Bayes model match the pattern exhibited by the logistic-regression Joint Culture model, except that the Naive-Bayes Joint-Culture model increases steadily as more groups are added with a maximum performance with 5 latent dimensions. With 5 latent dimension the gap between the two ML approaches shrinks to 2.8 percent (73.2 to 76.0).

On average the Joint Culture model shows a 1.8 percent improvement (74.3 to 76.1) over the culture neutral model for the L2-Logistic Regression, while it is a larger 2.2 percent for the Naive Bayes based approach (71.2 to 73.4). A comparison of the error reduction shows that the cultural integration is very promising. While the L2-loss logistic regression provides an 11 percent error reduction over the Naive-Bayes, the joint culture model achieves a comparable 7-9 percent reduction in error over the L2-loss regression and the Naive-Bayes model.

The improvements are strongest for positive self directed reward factor, skill factor, and lack of control factor. Interestingly, the models also exhibited considerable variation in accuracies across the different cultures, for example utilizing 3 dimensions positive rewards for others in one culture is recognized at 92 percent (this group contains 43,556 tweets), while for another culture of approximately the same size it is only recognized at 77 percent. Unfortunately, when moving to 4 dimension our clustering algorithm splits the group at 92 percent into two groups where the factor can only be recognized with an average of 88 percent accuracy. This suggests that more complex clusterings strategies within the latent space would be beneficial.

While not shown in the table for space reasons, we also tested the joint culture model utilizing a random assignment of authors to cultures, instead of relying on the assignment produced by the SVD-model. As expected the random model performed, on average, at approximately the same level as the general model, 74.6% compared to 74.5% respectively. Though the random culture model exhibited considerable variation in relation to the real joint culture model across the different factors. This evidence reinforces the idea that the latent cultures are coherent and that individuals within those cultures express the factors of intentionality in similar ways.

Table 4: Example cultures and the tags that are commonly associated with that factor.

Cultural Label	Cognitive Factor	Common Tags
Alterantive Medicine Health	Positive Rewards Self	#almond, #radish, #curd
Geek Interest	Positive Rewards Self	#theobroma, #freefiction, #nanotech
Teenagers	Positive Rewards Self	#bored, #me, #cute
Urban Hip/Hop	Positive Rewards Self	#bosslife, #teamfastfollow, #indiecharts
Martial Fitness	Goals	#healthynews, #fitnessimages, #fitso
Hip/Hop	Goals	#soundcloud, #support, #dl
General Religion	Goals	#singer, #jesus, #judas

Inspections of tweets where the cognitive factors have been discovered suggest that many times the hashtags 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 even though neither were identified by our annotators initially. These two tags are used by individuals on the first and second day of pursuing a goal.

The results presented in this section suggest that breaking down the authors by culture before learning models linking the hashtags marking expressions of the cognitive factors of intentionality to language provides a significant benefit. It also hints at some interesting differences between the groups. In the next section we briefly explore some of those differences.

5   Investigating Cultural Differences in the Language of Intentionality

We investigated the cultural discriminations made by the model by looking at the hashtags that were the most popular for each culture. Two annotators provided labels for each of the cultures based on the most frequent hashtags for that culture. We found that some of the cultures could easily be labeled based on their differential use of topical hashtags. Many of the latent cultures reflected notions of distinctions between cultural (or sub-cultural) groups, such as along political orientation or socio-demographics (urban, hipster, university students, single mothers, and political activist). In addition to the latent cultures that weighed on group identity, some of the other clusters captured more topical information, such as being fitness oriented or discussions focused around sex.

The cultural distinctions allowed us to quantify the differences in the event and intentionality associations across the cultures and differences in expressions indicating cognitive factors of intentionality. For instance, activists and urban individuals were most likely to produce tweets expressing control over situations. There were also groups, such as the camaraderie group where individuals typically set goals that will benefit a group in some way as well as the individual. In most of these cases, the author is the member of a team or some other group that will be engaging in a cooperative or competitive activity. Some authors from this cultural group express goals of providing direct or moral support to specific teams or groups of which they are not members. Others have goals of attending group events or gatherings with no particular membership. In most cases, goals in this culture are associated with positive rewards or defeating an opponent.

Table 4 shows the most probable tags by cognitive factor for some of the more interesting groups. These lists were generated by first eliminating all tags from the culture that were not predictive of the culture. To do this, we generated an estimate of the mean and variance for each hashtag in our dataset across all of the different cultures. We then eliminated all tags where the probability of the tag given the culture was not significantly different than its estimate given the general population. This has the effect of removing that hashtags that signaled the cognitive factors because they had a fairly general distribution across the cultures.

6   Conclusion

In this paper we presented a novel approach for identifying factors of intentionality in tweets. Further, we showed how a latent cultural model could be used to enhance those identifications through an improved understanding of how these factors are expressed across the various cultures. The latent cultural dimensions identified by the model correspond well with real cultural demographic information.

This work presents several exciting possibilities, while Twitter is notoriously difficult for traditional natural language processing work because it doesn't follow established syntactic and semantic conventions, models learned over Twitter data are able to transfer to other types of social media, such as user-generated content sites (Tomlinson et al., 2014a). Hashtags provide a very interesting form of distant annotation that could reduce the amount of time and effort required to create models which capture a nuanced understanding of social or psychological pragmatics,

such as social acts (Bender et al., 2011; Bracewell et al., 2012a), thus making the exploration of a richer language understanding more tractable.

Lastly, we have also shown that the models provide an ability to look at differences between cultures in the how and when of their expressions of factors relating to intentionality. People express lots of goals, but what affects when they actually intent them. These models should be able to provide a novel view on the pulse of a city (Rios and Lin, 2013) or citizens' cognitive responses to events (Dodds et al., 2011). We can use these techniques to identify what events make people establish new goals or instill feelings of a loss of control?

Acknowledgment

References

Icek Ajzen and Martin Fishbein. 1977. Attitude-behavior relations: A theoretical analysis and review of empirical research. *Psychological Bulletin*, 84(5):888–918.

Icek Ajzen. 1991. The Theory of Planned Behavior. *Organizational Behavior and Human Decision Processes*, 50:179–211.

D. Albarracin, J. Hepler, and M. Tannenbaum. 2011. General Action and Inaction Goals: Their Behavioral, Cognitive, and Affective Origins and Influences. *Current Directions in Psychological Science*, 20(2):119–123, April.

Nilanjan Banerjee, Dipanjan Chakraborty, Anupam Joshi, Sumit Mittal, Angshu Rai, and B. Ravindran. 2012. Towards Analyzing Micro-Blogs for Detection and Classification of Real-Time Intentions. *ICWSM*.

E.M. Bender, J.T. Morgan, Meghan Oxley, Mark Zachry, Brian Hutchinson, Alex Marin, Bin Zhang, and Mari Ostendorf. 2011. Annotating social acts: Authority claims and alignment moves in wikipedia talk pages. *ACL HLT 2011*, (June):48.

David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3:993–1022.

David B Bracewell, Marc T Tomlinson, Mary Brunson, Jesse Plymale, Jiajun Bracewell, and Daniel Boerger. 2012a. Annotation of Adversarial and Collegial Social Actions in Discourse. In *6th Linguistic Annotation Workshop*, number July, pages 184–192.

David B Bracewell, Marc T. Tomlinson, and Hui Wang. 2012b. Identification of Social Acts in Dialogue. In *COLING*, number December 2012, pages 375–390.

Xi Chen, Bing Bai, Qihang Lin, and Jaime G Carbonell. 2011. Sparse Latent Semantic Analysis. In *SDM*.

Zhiyuan Chen, Bing Liu, Meichun Hsu, Malu Castellanos, and Riddhiman Ghosh. 2013. Identifying Intention Posts in Discussion Forums. In *NAACL-HLT*, number June, pages 1041–1050.

Raviv Cohen and Derek Ruths. 2013. Classifying Political Orientation on Twitter : It s Not Easy ! In *ICWSM-2013*, pages 91–99.

Dmitry Davidov, Oren Tsur, and Ari Rappaport. 2010. Enhanced Sentiment Learning Using Twitter Hashtags and Smileys. In *Coling*, number August, pages 241–249.

PS Dodds, KD Harris, and IM Kloumann. 2011. Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter. *PloS one*, 6.

David Dowty. 1991. Thematic Proto-Roles and Argument Selection. *Linguistic Society of America*, 67(3):547–619.

RE Fan, KW Chang, CJ Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. 2008. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9(2008):1871–1874.

Stacey R. Finkelstein and Ayelet Fishbach. 2012. Tell Me What I Did Wrong: Experts Seek and Respond to Negative Feedback. *Journal of Consumer Research*, 39(1):22–38, June.

Kevin Gimpel, Nathan Schneider, Brendan O Connor, Dipanjan Das, Daniel Mills, Jacob Eisenstein, Michael Heilman, Dani Yogatama, Jeffrey Flanigan, and Noah A Smith. 2011. Part-of-Speech Tagging for Twitter : Annotation , Features , and Experiments. In *Proceedings of the Association for Computational Linguistics*, number 2.

B Grosz and C Sidner. 1986. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3).

Geert Hofstede. 1980. *Culture's consequences: International differences in work-related values*. Sage Publications, Inc.

R. Kannan, S. Vempala, and A. Vetta. 2004. On clusterings: Good, bad and spectral. *Journal of the ACM (JACM)*, 51(3):497–515.

J. Knobe. 2003. Intentional action and side effects in ordinary language. *Analysis*, 63(3):190–194, July.

E. A. Locke. 1968. Toward a theory of task motivation and incentives. *Organizational behavior and human performance*, 3(2).

Bertram Malle and J. Knobe. 1997. The Folk Concept of Intentionality. *Journal of Experimental Psychology*, 33(2):101–121.

Marco Pennacchiotti and Ana-maria Popescu. 2011. to Twitter User Classification. In *ICWSM'11*, pages 281–288.

Miguel Rios and Jimmy Lin. 2013. Visualizing the" Pulse" of World Cities on Twitter. *Seventh International AAAI Conference on Weblogs . . .*, pages 717–720.

Alan Ritter, Colin Cherry, and Bill Dolan. 2009. Unsupervised Modeling of Twitter Conversations. In *HTL-NAACL*.

Beaux Sharifi, Mark-anthony Hutton, and Jugal Kalita. 2010. Summarizing Microblogs Automatically. In *ACL-HLT*, number June, pages 685–688.

Steven a. Sloman, Philip M. Fernbach, and Scott Ewing. 2012. A Causal Model of Intentionality Judgment. *Mind & Language*, 27(2):154–180, April.

Michael Tomasello, Malinda Carpenter, Josep Call, Tanya Behne, and Henrike Moll. 2005. Understanding and sharing intentions: the origins of cultural cognition. *The Behavioral and brain sciences*, 28(5):675–91; discussion 691–735, October.

Marc T Tomlinson, David Bracewell, Wayne Krug, and David Hinote. 2014a. # impressme : The Language of Motivation in User Generated Content. In *CICLING*.

Marc T Tomlinson, David Bracewell, Wayne Krug, David Hinote, and Mary Draper. 2014b. # mygoal : Finding Motivations on Twitter. In *LREC - 2014*. ELRA.

Faiyaz Al Zamal, Wendy Liu, and Derek Ruths. 2012. Homophily and Latent Attribute Inference : Inferring Latent Attributes of Twitter Users from Neighbors. In *ICWSM-2012*.