Predicting Attitude and Actions of Twitter Users

Jalal Mahmud[1], Geli Fei[2*], Anbang Xu[1], Aditya Pal[3*], Michelle Zhou[4*]

[1]IBM Research - Almaden, jumahmud@us.ibm.com, anbangxu@us.ibm.com
[2]University of Illinois at Chicago, gfei2@uic.edu
[3]Facebook Inc, aditya.pal@gmail.com
[4]Juji Inc, mzhou@acm.org

ABSTRACT

In this paper, we present computational models to predict Twitter users' attitude towards a specific brand through their personal and social characteristics. We also predict their likelihood of taking different actions based on their attitudes. In order to operationalize our research on users' attitude and actions, we collected ground-truth data through surveys of Twitter users. We have conducted experiments using two real world datasets to validate the effectiveness of our attitude and action prediction framework. Finally, we show how our models can be integrated with a visual analytics system for customer intervention.

Author Keywords

Attitude; Social Media; Twitter; Brand

ACM Classification Keywords

H.5.2 [Information Interfaces and Presentation]: User Interfaces - Interaction styles.

General Terms

Human Factors; Design; Measurement.

INTRODUCTION

In the last few years, social media such as Twitter has emerged, and different brands have social media presence to attract their potential customers. People express various opinions about such brands in social media. Some people may like a brand (e.g., Delta Airlines, Fitbit), some may show neutral attitude, and others may dislike the brand. Some may have formed an attitude towards a brand very recently, and others may have an attitude for quite a long time. Some people have attitude with higher confidence than others, some may remember their attitude well and some are more likely to change their attitudes. People also take different actions (e.g., buy a product/service corresponding to that brand, recommend others to buy) based on their attitude about a brand. Prior works on sentiment/opinion analysis [3, 4, 6, 7, 10, 11, 12, 14] can be useful to know whether a user may like/dislike a brand. A recent work on attitude modeling [1] also describes inferring attitude towards controversial topics in terms of sentiment, opinion and actions. However, such works do not address whether such attitudes are persistent (e.g., an individual formed an attitude for a long time) or temporary (e.g., an individual formed an attitude recently). They also do not provide the strength of attitude (e.g., whether the individual has attitude with high/low confidence). Furthermore, they do not address how well a user remembers his/her attitude or whether the user is likely to change the attitude. Such fine-grained information of consumer attitude can be useful for social media marketers who would directly engage such consumers on social media platforms [5]. Furthermore, predicting social media actions (e.g., retweeting a tweet) as described in Gao et al. is inadequate for such scenario where marketers would be more interested to know whether a consumer will take actions outside the social media.

Motivated by such a need, we present computational models to predict a Twitter user's attitude in terms of a number of characteristics such as *attitude favorability* (How much a consumer likes or dislikes an attitude object), *attitude persistence* (whether an attitude is persistent), *attitude confidence* (strength of attitude), *attitude accessibility* (How well a consumer remembers attitude about the object) and *attitude resistance* (How likely a consumer keeps the present attitude). Our work is inspired by marketing literature where attitude is described in terms of such characteristics [13].

Since there are no publicly available ground truth data of attitude characteristics, we have collected such data using self-report surveys conducted among Twitter users. Using the ground-truth data, we developed statistical models to predict users' attitude. Our classification based models of attitude characteristics are based on features extracted from users' historical tweets. Such models can classify whether a user has specific characteristics of attitude, and also output the likelihood of those attitude characteristics. We have also developed statistical models to infer likelihood of different action intention (e.g., intention to buy a product) based on one's attitude. Similar to our models for predicting different attitude characteristics, our models for predicting action

* This work was done while the author was with IBM Research.

IUI '16, March 7–10, 2016, Sonoma, CA, USA

DOI: http://dx.doi.org/10.1145/2856767.2856800

intentions are also trained from features derived from users' historical tweets.

We performed extensive experiments using two real world datasets to validate the effectiveness of our models. For attitude characteristics, we observed mixed result. While attitude favorability can be predicted within 65-69% AUC, prediction accuracies of other characteristics are low to moderate (52-59% AUC). Action intentions can be predicted within 56-67% AUC. We have also integrated our prediction models with a visual analytics system that recommends Twitter users with specific attitude towards a brand, and thus allows potential intervention. Below we list the summary of contributions of this work:

- A survey study of understanding attitude of Twitter users towards multiple brands.
- Models to predict attitude of Twitter users towards a brand in terms of a set of characteristics.
- Models to predict users' intention to take different actions based on their attitudes.
- Experiments demonstrating the effectiveness of our models.
- A visual analytic system that integrates our models and allows customer intervention.

SURVEY STUDY - METHODOLOGY

We collected two real-world datasets from Twitter. Our first dataset is about "Delta" airlines, and the second dataset is about "Fitbit" exercise equipment.

Survey Questions: We created the surveys using SurveyGizmo (http://www.surveygizmo.com/), a popular survey-building tool. Our survey questions were designed to capture various aspects of attitude [13]. Table 1 and 2 shows questions that were included in the Delta survey. Similarly, Table 3 and 4 shows questions that were included in the Fitbit survey. Response to each question was on 5-point Likert scale. Note that, some attitude variables are measured using multiple questions. Response to such variables is computed as an average of the response of those questions. A high value of favorability means user likes the brand more, and low value means the opposite. A high/low value of persistence means user has more/less persistence attitude. Similarly, a high value of accessibility means user can remember the attitude easily, and low value means the opposite. A high value of resistance means user is more likely to stay with the brand, and low value means the opposite. A low value on the response to each action intention means the user is less likely to perform the action, and high value means the opposite.

Survey Participants:

We identified 7534 Twitter users who tweeted about "Delta" airlines during March 2014 to June 2014. They were identified using the keyword @delta in Twitter's Search API. We sent them requests to participate in our survey. To do so, we constructed 6 Twitter accounts and these accounts were used for sending the surveys. These accounts were constructed in a way to appear as genuine as possible, so that a survey request from them does not appear to be a phishing or marketing campaign to target users and the risk of them getting marked as spam by Twitter is mitigated. We offered $50 Amazon gift card to 1 out of every 100 survey participants. Similarly, we identified 5261 Twitter users who tweeted about "Fitbit" exercise equipment (identified using the keyword @fitbit in Twitter's Search API), and sent request to participate in our survey. Similar to our "Delta" airlines survey, we created 6 Twitter accounts to send survey requests to Twitter users and we offered $50 Amazon gift card to 1 out of every 100 survey participants. We ensured that each participant took the survey only once. We also made sure study participants' privacy was protected. We anonymized our ground truth data so that participants were not identifiable afterwards. Before taking the survey, each participant was told that their survey responses will be anonymized, and they gave us explicit consent (by reading consent statement and indicating in the user interface of the survey) to pull their tweets from Twitter. We did not ask their Twitter id to protect their privacy, instead we asked them to give access to their tweets using oAuth. These tweets were linked to their survey responses but was anonymized, so users ID or personal data was not stored.

Attitude Characteristics	Questions
Favorability	How much have you liked your travel experience with Delta Airlines?
Persistence	How often have you used Delta Airlines for your travel? How long have you used Delta Airlines for your travel?
Confidence	Based on your answers, how certain are you about your answers?
Accessibility	How well do you remember your attitude about Delta Airlines?
Resistance	How likely will you switch to another airlines if Delta reduces efficiency of service? How likely will you switch to another airlines if Delta reduces comfort of service? How likely will you switch to another airlines if Delta increases cost of service?

Table 1. Questions to assess attitude about Delta

Survey Responses:

823 users responded to our Delta survey (10.9% response rate) and 507 users responded to our Fitbit (9.6% response rate) survey. We manually inspected survey responses and removed incomplete and inconsistent responses. Finally, we had 751 survey responses for Delta and 447 survey

responses for Fitbit. For each user who responded, we collected their most recent 3200 tweets (max limit enforced by Twitter) using Twitter's REST API. If they had less than 3200 tweets, then we collected all their tweets. We also checked their tweets and they were all English.

Action Intentions	Questions
Buy	How likely are you going to buy ticket of your next trip from Delta airlines?
Recommend	How likely are you going to recommend others to fly by Delta airlines?
Prohibit	How likely are you going to tell others not to fly by Delta airlines?

Table 2. Questions in Delta survey to assess action intentions

Attitude Characteristics	Questions
Favorability	How much have you liked your experience with Fitbit?
Persistence	How long have you used Fitbit?
Confidence	Based on your answers, how certain are you about your answers?
Accessibility	How well do you remember your attitude about Fitbit?
Resistance	How likely will you switch to another fitness device if Fitbit reduces efficiency (e.g., calorie tracking)? How likely will you switch to another fitness device if Fitbit reduces comfort (e.g., comfort to wear)? How likely will you switch to another fitness device if Fitbit increases cost? How likely will you switch to another fitness device if Fitbit reduces visual attractiveness?

Table 3. Questions to assess attitude about Fitbit

Action Intentions	Questions
Buy	How likely are you going to buy next fitness device from Fitbit?
Recommend	How likely are you going to recommend others to buy fitness device from Fitbit?
Prohibit	How likely are you going to tell others not to use Fitbit?

Table 4. Questions in Fitbit survey to assess action intentions

CLASSIFICATION APPROACHES

We developed statistical models to classify each attitude characteristics and action intention. Such models used a set of features extracted from users' historical tweets. For simplicity, we developed binary classifiers for each attitude characteristics. Thus, we first converted each attitude characteristics into binary values (using mean as threshold).

Feature Extraction.

- *Unigram features:* This feature represents all unigrams extracted from tweets.
- *Sentiment features*: We use a sentiment/opinion dictionary that contains a list of words with their positive/negative sentiment polarity. We count total number of positive/negative words in user's tweets and used that positive and negative counts as features.
- *Context-based Sentiment/Opinion feature*: This is similar to the above sentiment/opinion feature; however, it looks for sentiment words that appear in the surrounding area of the brand name (e.g., textual patterns like "awesome delta"). Thus, we counted how many times positive sentiment words in the dictionary co-occur with the brand name and how many times negative sentiment words in the dictionary co-occur with the brand name and used those counts as feature values.
- *Domain-specific sentiment feature*: This feature is similar to the above *Context-based Sentiment/Opinion feature*, however, it is computed by matching words in users' tweets with a domain-specific sentiment dictionary which we constructed from training users' tweets.
- Length of use feature: This feature captures the attitude-persistence of a user and is obtained by taking the timestamp difference of a user's latest and oldest mention of the brand.
- *Frequency feature*: This feature represents how often the user mentions the brand.

Statistical Models.

Once we computed the above mentioned features, we developed statistical models using WEKA [15]. We tried a number of classifiers such as Naive Bayes, SMO (SVM), Random Forest from WEKA and performed 5-fold cross validation. SMO and Random Forest based classifier achieved comparable performance. In experiment section, we report experimental result for SMO classifier.

EXPERIMENTS

Here we describe experiments we have performed to validate the effectiveness of our approach.

Attitude Classification

Table 5 shows the result of our attitude prediction for Delta and Fitbit datasets in terms of F1 and AUC. We observe mixed result for attitude characteristics prediction. Attitude favorability can be predicted with reasonable accuracy (65%-69% AUC). However for other attitude characteristics, we found moderate to low prediction accuracy. It is more intuitive that textual and sentiment based feature extracted from users' historical tweets contain predictive information to predict attitude favorability (i.e., how much they like/dislike a brand). Attitude accessibility and resistance were more difficult to predict which could be either due to the reason that our survey users did not provide very reliable response for those dimensions or users' historical tweets did not contain enough predictive feature for predicting them.

	F1		ROC Area (AUC)	
	Delta	Fitbit	Delta	Fitbit
Favorability	0.71	0.68	0.69	0.65
Persistence	0.61	0.57	0.59	0.57
Confidence	0.58	0.56	0.56	0.55
Accessibility	0.55	0.53	0.53	0.52
Resistance	0.53	0.53	0.52	0.52

Table 5. Result of Attitude Classification

	F1		ROC Area (AUC)	
	Delta	Fitbit	Delta	Fitbit
Buy	0.68	0.65	0.67	0.63
Recommend	0.65	0.60	0.64	0.58
Prohibit	0.60	0.58	0.58	0.56

Table 6. Result of Action Intention Classification

Action Intention Classification

Table 6 shows the result of action intention prediction for both datasets in terms of F1 and AUC. Overall, action intentions can be predicted within reasonable accuracy (56% to 67% AUC). In each case, buy action achieves best performance, while prohibit action seems harder to predict. We think one of the difficulties of predicting prohibit action comes from the particularly unbalanced training data, as prohibit action is an extreme action that not many users would do.

SYSTEM FOR CUSTOMER INTERVENTION

We integrated our prediction models with a visual analytics system for customer intervention. We iteratively designed such a system. First, the system uses simple keyword filtering to identify a set of Twitter users who have recently discussed the brand in their tweets. Their attitudes and action intentions are computed based on their tweets and presented to agents for intervention. The design is shown in Figure 1 and includes two main components: the *attitude & intention component* and the *detailed view component*. The *attitude & intention component* (right in Figure 1) shows an overview of customer attitudes and action intentions in a number of dimensions including favorability, persistence, confidence, accessibility, resistance, buy, recommend, and prohibit. A bar chart visualizes the distribution of the customers' attitude or action intention values along a dimension. The number in a bar denotes the total number of customers in the corresponding segment (see Figure 1). The system allows agents to create visual filters on any individual bar chart in the *attitude & intention component* by selecting a range on the axis, and customers within the selected range would then be shown in the *detailed view component* (left in Figure 1) to reflect the selection. The filtering feature helps agents identify target customers for intervention.

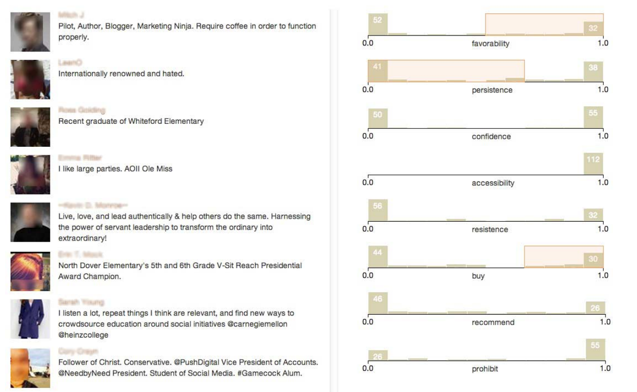

Figure 1. A screenshot of the visual analytics system

LIMITATIONS

Our work has several limitations. We found that accuracy of our prediction models are not high. This can be due to various factors. First, our survey responders might not provide reliable response. We understand that self reported ground truth is not perfect. In our experiments, we did manually check responses of the survey to see if there was any inconsistency in response, and deleted some inconsistence response. In future we will explore other forms of ground truth collection. Second, our models are trained such that features extracted from historical tweets of users are used to predict their attitude and action intention. Such tweets may not contain enough information to predict users' attitude and intentions. In future, we plan to experiment with more features (e.g., features used in [8] to predict re-tweeters) and their relationships to find if prediction accuracy can be improved.

CONCLUSIONS AND FUTURE WORK

This work is a first exploration to predict one's attitude towards a brand in terms of a set of characteristics, and likelihood to take different actions based on attitudes. Our models are trained and tested using two real-world datasets. We have found that most of the attitude characteristics and action intentions can be predicted with moderate accuracy (60-65%). Furthermore, we have integrated our prediction models to a visualization interface to demonstrate usage in customer intervention. In future, we plan to explore how to construct attitude models which are easily scalable across multiple brands and apply our models in real-world scenarios.

REFERENCES

1. Gao, H. Mahmud, J., Chen, J., Nichols, J, Zhou, M.X. Modeling User Attitude Toward Controversial Topics in Online Social Media. In *Proc.* ICWSM 2014.
2. Hoyer, W.D., Maclnnis, D.J. Pieters, R. Consumer Behavior. 5th Edition, 2008.
3. Hu, X.; Tang, J.; Gao, H.; and Liu, H. 2013a. Unsupervised sentiment analysis with emotional signals. In Proc. of the WWW 2013, 607–618.
4. Hu, X.; Tang, L.; Tang, J.; and Liu, H. 2013b. Exploiting social relations for sentiment analysis in microblogging. In Proc. of the WSDM 2013, 537–546.
5. Jansen, B. J., Zhang, M., Sobel, K., et al. 2009. Twitter power: Tweets as electronic word of mouth. Journal of American Society for Information Science & Technology, 60(11), 2169-2188.
6. Jiang, L.; Yu, M.; Zhou, M.; Liu, X.; and Zhao, T. 2011. Target dependent twitter sentiment classification. In Proc. of the HLT '11, 151–160.
7. Kim, J.; Yoo, J.; Lim, H.; Qiu, H.; Kozareva, Z.; and Galstyan, A. 2013. Sentiment prediction using collaborative filtering. In Proc. of the ICWSM 2013.
8. Lee, K.; Mahmud, J.; Chen, J., Zhou, M.X., Nichols, J. 2014. Who Will Retweet This? Automatically Identifying and Engaging Strangers on Twiteer to Spread Information. In Proc. IUI 2014.
9. Li, F.; Huang, M.; and Zhu, X. 2010. Sentiment analysis with global topics and local dependency. In *Proc.* AAAI 2010.
10. Li, D.; Shuai, X.; Sun, G.; Tang, J.; Ding, Y.; and Luo, Z. 2012. Mining topic-level opinion influence in microblog. In *Proc.* CIKM 2012.
11. Lin, C., and He, Y. 2009. Joint sentiment/topic model for sentiment analysis. In *Proc.* CIKM 2009.
12. Pang, B. and Lee, L. 2008. Opinion mining and sentiment analysis. Foundations and Trends in Information Retrieval (2).
13. Schiman, L., and Kanuk, L. 2010. Consumer Behavior, 10th edition. Prentice Hall.
14. Tan, C.; Lee, L.; Tang, J.; Jiang, L.; Zhou, M.; and Li, P. 2011. User-level sentiment analysis incorporating social networks. In Proc. of KDD 2011.
15. Witten, I.H., Frank, E., and Hall, M.A. 2011. Data mining: Practical machine learning tools and techniques, 3rd Edition. Morgan Kaufmann.