Using Link Analysis to Discover Interesting Messages Spread Across Twitter

Min-Chul Yang[†] and **Jung-Tae Lee**[‡] and **Hae-Chang Rim**[†]

[†]Dept. of Computer & Radio Communications Engineering, Korea University, Seoul, Korea
[‡]Research Institute of Computer Information & Communication, Korea University, Seoul, Korea
{mcyang,jtlee,rim}@nlp.korea.ac.kr

Abstract

Twitter, a popular social networking service, enables its users to not only send messages but re-broadcast or *retweet* a message from another Twitter user to their own followers. Considering the number of times that a message is retweeted across Twitter is a straightforward way to estimate how interesting it is. However, a considerable number of messages in Twitter with high retweet counts are actually mundane posts by celebrities that are of interest to themselves and possibly their followers. In this paper, we leverage retweets as implicit relationships between Twitter users and messages and address the problem of automatically finding messages in Twitter that may be of potential interest to a wide audience by using link analysis methods that look at more than just the sheer number of retweets. Experimental results on real world data demonstrate that the proposed method can achieve better performance than several baseline methods.

1 Introduction

Twitter (http://twitter.com) is a popular social networking and microblogging service that enables its users to share their status updates, news, observations, and findings in real-time by posting text-based messages of up to 140 characters, called *tweets*. The service rapidly gained worldwide popularity as a communication tool, with millions of users generating millions of tweets per day. Although many of those tweets contain valuable information that is of interest to many people, many others are mundane tweets, such as *"Thanks guys for the birthday wishes!!"* that are of interest only to the authors and users who subscribed to their tweets, known as *followers*. Finding tweets that are of potential interest to a wide audience from large volume of tweets being accumulated in real-time is a crucial but challenging task. One straightforward way is to rely on the numbers of times each tweet has been propagated or *retweeted* by readers of the tweet. Hong et al. (2011) propose to regard retweet count as a measure of popularity and present classifiers for predicting whether and how often new tweets will be retweeted in the future. However, mundane tweets by highly popular users, such as celebrities with huge numbers of followers, can record high retweet counts. Alonso et al. (2010) use crowdsourcing to categorize a set of tweets as "only interesting to author and friends" and "possibly interesting to others" and report that the presence of a URL link is a single, highly effective feature for distinguishing interesting tweets with more than 80% accuracy. This simple rule, however, may incorrectly recognize many interesting tweets as not interesting, simply because they do not contain links. Lauw et al. (2010) suggest several features for identifying interesting tweets but do not experimentally validate them.

In this study, we follow the definition of interesting tweets provided by Alonso et al. (2010) and focus on automatic methods for finding tweets that may be of potential interest to not only the authors and their followers but a wider audience. Since retweets are intended to spread tweets to new audiences, they are often a recommendation or, according to Boyd et al. (2010), productive communication tool. Thus, we model Twitter as a graph con-

sisting of user and tweet nodes implicitly connected by retweet links, each of which is formed when one user retweets what another user tweeted. We present a variant of the popular HITS algorithm (Kleinberg, 1999) that exploits the retweet link structure as an indicator of how interesting an individual tweet is. Specifically, we draw attention on the fact that not all retweets are meaningful. Some users retweet a message, not because of its content, but only because they were asked to, or because they regard retweeting as an act of friendship, loyalty, or homage towards the person who originally tweeted (Boyd et al., 2010). The algorithm proposed in this paper is designed upon the premise that not all retweet links are created equal, assuming that some retweets may carry more importance or weight than others. Welch et al. (2011) and Romero et al. (2011) similarly extend link analysis to Twitter, but address essentially different problems. We conduct experiments on real world tweet data and demonstrate that our method achieves better performance than the simple retweet count approach and a similar recent work on Twitter messages (Castillo et al., 2011) that uses supervised learning with a broad spectrum of features.

2 Proposed Method

We treat the problem of finding interesting tweets as a ranking problem where the goal is to derive a scoring function which gives higher scores to interesting tweets than to uninteresting ones in a given set of tweets. To derive the scoring function, we adopt a variant of HITS, a popular link analysis method that emphasizes mutual reinforcement between authority and hub nodes (Kleinberg, 1999).

Formally, we model the Twitter structure as directed graph $G = (N, E)$ with nodes N and directional edges E. We consider both users $U = \{u_1, \ldots, u_{n_u}\}$ and tweets $T = \{t_1, \ldots, t_{n_t}\}$ as nodes and the retweet relations between these nodes as directional edges. For instance, if tweet t_a, created by user u_a, retweets t_b, written by user u_b, we create a retweet edge e_{t_a, t_b} from t_a to t_b and another retweet edge e_{u_a, u_b} from u_a to u_b.[1] Strictly speaking, G has two subgraphs, one based only on the user nodes and another based on the tweet nodes. Instead of running HITS on the tweet subgraph right away,

we first run it on the user subgraph and let tweets inherit the scores from their publishers. Our premise is that the scores of a user is an important prior information to infer the scores of the tweets that the user published.

User-level procedure: We first run the algorithm on the user subgraph. $\forall u_i$, we update the authority scores $A(u_i)$ as:

$$\sum_{\forall j : e_{u_j, u_i} \in E} \frac{|\{u_k \in U : e_{u_j, u_k} \in E\}|}{|\{k : e_{u_j, u_k} \in E\}|} \times H(u_j) \quad (1)$$

Then, $\forall u_i$, we update the hub scores $H(u_i)$ to be:

$$\sum_{\forall j : e_{u_i, u_j} \in E} \frac{|\{u_k \in U : e_{u_k, u_j} \in E\}|}{|\{k : e_{u_k, u_j} \in E\}|} \times A(u_j) \quad (2)$$

A series of iterations is performed until the scores are converged. After each iteration, the authority/hub scores are normalized by dividing each of them by the square root of the sum of the squares of all authority/hub values. When this user-level stage ends, the algorithm outputs a function $S_{U_A} : U \rightarrow [0, 1]$, which represents the user's final authority score, and another function $S_{U_H} : U \rightarrow [0, 1]$, which outputs the user's final hub score. Note that, unlike the standard HITS, the authority/hub scores are influenced by edge weights that reflect the retweet behaviors of individual users. The idea here is to dampen the influence of users who devote most of their retweet activities toward a very few other users, such as celebrities, and increase the weight of users who retweet many different users' tweets. To demonstrate the effectiveness of the parameter, we have done some preliminary experiments. The column User$_{from}$ in Table 1 shows the retweet behavior of users who retweeted tweets belonging to "uninteresting" and "interesting" classes observed in our Twitter dataset. The values are calculated by the ratio of all other users that a user retweeted to all retweet outlinks from the user; a value closer to 1 means that outlinks are pointed to many different users.[2] We observe that the value for users who retweeted interesting tweets is shown to be higher, which means that they tend to retweet messages from many different users, more than users who retweeted uninteresting ones.

---

[1] Note that two user nodes can have multiple edges.

[2] For calculating the ratios, we limit the target to users who retweeted two or more times in our dataset.

Class	User$_{from}$	$F = \alpha$	#
Not Interesting	0.591	0.252	1985
Possibly Interesting	0.711	0.515	1115
Both	0.634	0.346	3100

Table 1: Dataset analysis.

Tweet-level procedure: After the user-level stage, we start computing the scores of the tweet nodes. In each iteration, we start out with each tweet node initially inheriting the scores of its publisher. Let $P : T \rightarrow U$ be a function that returns the publisher of a given tweet. $\forall t_i$, we update $A(t_i)$ to be:

$$S_{U_A}(P(t_i)) + \sum_{\forall j: e_{t_j, t_i} \in E} F(e_{t_j, t_i}) \times H(t_j) \quad (3)$$

Then, $\forall t_i$, we update $H(t_i)$ to be:

$$S_{U_H}(P(t_i)) + \sum_{\forall j: e_{t_i, t_j} \in E} F(e_{t_i, t_j}) \times A(t_j) \quad (4)$$

where $F(e_{t_a, t_b})$ is a parameter function that returns $\alpha > 1$ if $P(t_a)$ is not a follower of $P(t_b)$ and 1 otherwise. It is intuitive that if users retweet other users' tweets even if they are not friends, then it is more likely that those tweets are interesting. The column $F = \alpha$ in Table 1 shows the ratio of all unfollowers who retweeted messages in a particular class to all users who retweeted messages in that class, observed in our dataset. We observe that users retweet interesting messages more, even when they do not follow the publishers. Similar observation has also been made by Recuero et al. (2011). After each iteration, the authority/hub scores are normalized as done in the user-level. After performing several iterations until convergence, the algorithm finally outputs a scoring function $S_{T_A} : T \rightarrow [0, 1]$, which represents the tweet node's final authority score. We use this function to produce the final ranking of tweets.

Text pattern rules: We observe that in some cases users retweet messages from their friends, not because of the contents, but via retweet requests to simply evoke attention. To prevent useless tweets containing such requests from receiving high authority scores, we collect 20 simple text pattern matching rules that frequently appear in those tweets. Specifically, we let the rules make influence while updating the scores of tweets by modifying the summations in Eq. (3) and (4) respectively as:

$$\sum_{\forall j: e_{t_j, t_i} \in E} F(e_{t_j, t_i}) \times R(t_i) \times H(t_j) \quad (5)$$

$$\sum_{\forall j: e_{t_i, t_j} \in E} F(e_{t_i, t_j}) \times R(t_j) \times A(t_j) \quad (6)$$

where $R(t)$ is a rule-based function that returns 0 if tweet t contains one of the pre-defined text patterns and 1 otherwise. Such patterns include *"RT this if"* and *"If this tweet gets RT * times I will"*.

3 Experiment and Discussion

Our Twitter dataset is collected during 31 days of October 2011, containing 64,107,169 tweets and 2,824,365 users. For evaluation, we generated 31 immediate Twitter graphs composed of 1.5 million retweet links in average and 31 initially ranked lists of tweets, each consisting of top 100 tweets created on a specific date of the month with highest retweet counts accumulated during the next 7 days. Two annotators were instructed to categorize each tweet as interesting or not, by inspecting its content as done in the work of Alonso et al. (2010). In case of disagreement (about 15% of all cases), a final judgment was made by consensus between the two annotators. We observe that the ratio of tweets judged to be interesting is about 36%; the column '#' in Table 1 shows the actual counts of each class. The goal of this evaluation is to demonstrate that our method is able to produce better ranked lists of tweets by re-ranking interesting tweets highly.

Table 2 reports the ranking performance of various methods in terms of Precisions @10 and @20, R-Precision, and MAP. We compare our approach to four baselines. The first baseline, #RT, is obviously based on retweet counts; tweets with higher retweet counts are ranked higher. The second baseline, #URL+#RT, favors tweets that contain URL links (Alonso et al., 2010). Since it is less likely for a tweet to contain more than one link, we additionally use #RT to break ties in tweet ranking. Thirdly, HITS$_{original}$, is the standard HITS algorithm run on both user and tweet subgraphs that calculates authority/hub scores of a node purely by the sum of hub values that point to it and the sum of authority values that it points to, respectively, during iterations;

Method	P@10	P@20	R-Prec	MAP
#RT	0.294	0.313	0.311	0.355
#URL+#RT	0.245	0.334	0.362	0.361
$\text{HITS}_{original}$	0.203	0.387	0.478	0.465
$\text{ML}_{message}$	0.671	0.645	0.610	0.642
ML_{all}	0.819	0.795	0.698	0.763
$\text{HITS}_{proposed}$	**0.881**	**0.829**	**0.744**	**0.807**

Table 2: Performance of individual methods

Method	P@10	P@20	R-Prec	MAP
$\text{HITS}_{proposed}$	**0.881**	**0.829**	**0.744**	**0.807**
w/o User	0.677	0.677	0.559	0.591
w/o Tweet	0.861	0.779	0.702	0.772
w/o Rule	0.858	0.81	0.733	0.781

Table 3: Contributions of individual stages.

no other influential factors are considered in the calculations. Lastly, we choose one recent work by Castillo et al. (2011) that addresses a related problem to ours, which aims at learning to classify tweets as credible or not credible. Although interestingness and credibility are two distinct concepts, the work presents a wide range of features that may be applied for assessing interestingness of tweets using machine learning. For re-implementation, we train a binary SVM classifier using features proposed by Castillo et al. (2011), which include features from users' tweet and retweet behavior, the text of the tweets, and citations to external sources; we use the probability estimates of the learned classifier for re-ranking.[3] We use leave-one-out cross validation in order to evaluate this last approach, denoted as ML_{all}. $\text{ML}_{message}$ is a variant that relies only on message-based features of tweets. Our method, with α empirically set to 7, is denoted as $\text{HITS}_{proposed}$.

We observe that #RT alone is not sufficient measure for discovering interesting tweets. Additionally leveraging #URL helps, but the improvements are only marginal. By manually inspecting tweets with both high retweet counts and links, it is revealed that many of them were tweets from celebrities with links to their self-portraits photographed in their daily lives, which may be of interest to their own followers only. $\text{HITS}_{original}$ performs better than both #RT and #URL across most evaluation metrics but generally does not demonstrate good performance. $\text{ML}_{message}$ always outperform the first three significantly; we observe that tweet lengths in characters and in words are the two most effective message-based features for finding interesting tweets. The results of ML_{all} demonstrates that more

reasonable performance can be achieved when user- and propagation-based features are combined with message-based features. The proposed method significantly outperforms all the baselines. This is a significant result in that our method is an unsupervised approach that relies on a few number of tweet features and does not require complex training.

We lastly report the contribution of individual procedures in our algorithm in Table 3 by ablating each of the stages at a time. "*w/o User*" is when tweet nodes do not initially inherit the scores of their publishers. "*w/o Tweet*" is when tweets are re-ranked according to the authority scores of their publishers. "*w/o Rule*" is when we use Eq. (3) and (4) instead of Eq. (5) and (6) for updating tweet scores. We observe that the user-level procedure plays the most crucial role. We believe this is because of the ability of HITS to distinguish good "hub-users". Since authoritative users can post ordinary status updates occasionally in Twitter, we cannot always expect them to create interesting content every time they tweet. However, good hub-users[4] tend to continuously spot and retweet interesting messages; thus, we can expect the tweets they share to be interesting steadily. The role of hubs is not as revealed on the tweet side of the Twitter graph, since each tweet node can only have at most one retweet outlink. The exclusion of text pattern rules does not harm the overall performance much. We suspect this is because of the small number of rules and expect more improvement if we add more effective rules.

Acknowledgments

---

[3]We do not use some topic-based features in (Castillo et al., 2011) since such information is not available in our case.

[4]Often referred to as *content curators* (Bhargava, 2009).

References

Omar Alonso, Chad Carson, David Gerster, Xiang Ji, and Shubha U. Nabar. 2010. Detecting uninteresting content in text streams. In *Proceedings of the SIGIR 2010 Workshop on Crowdsourcing for Search Evaluation*, CSE '10, pages 39–42.

Rohit Bhargava. 2009. Manifesto for the content curator: The next big social media job of the future? `http://rohitbhargava.typepad.com/`.

Danah Boyd, Scott Golder, and Gilad Lotan. 2010. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In *Proceedings of the 2010 43rd Hawaii International Conference on System Sciences*, HICSS '10, pages 1–10.

Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information credibility on twitter. In *Proceedings of the 20th international conference on World wide web*, WWW '11, pages 675–684.

Liangjie Hong, Ovidiu Dan, and Brian D. Davison. 2011. Predicting popular messages in twitter. In *Proceedings of the 20th international conference companion on World wide web*, WWW '11, pages 57–58.

Jon M. Kleinberg. 1999. Authoritative sources in a hyperlinked environment. *J. ACM*, 46(5):604–632.

Hady W. Lauw, Alexandros Ntoulas, and Krishnaram Kenthapadi. 2010. Estimating the quality of postings in the real-time web. In *Proceedings of the WSDM 2010 Workshop on Search in Social Media*, SSM '10.

Raquel Recuero, Ricardo Araujo, and Gabriela Zago. 2011. How does social capital affect retweets? In *Proceedings of the 5th International AAAI Conference on Weblogs and Social Media*, ICWSM '11, pages 305–312.

Daniel M. Romero, Wojciech Galuba, Sitaram Asur, and Bernardo A. Huberman. 2011. Influence and passivity in social media. In *Proceedings of the 2011 European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, ECML-PKDD '11, pages 18–33.

Michael J. Welch, Uri Schonfeld, Dan He, and Junghoo Cho. 2011. Topical semantics of twitter links. In *Proceedings of the 4th International Conference on Web Search and Web Data Mining*, WSDM '11, pages 327–336.