
Characterizing the Effectiveness of Twitter Hashtags to Detect and Track Online Population Sentiment

Glívia A. R. Barbosa

Depart. of Computer Science
Federal University of Minas
Gerais
gliviaangelica@dcc.ufmg.br

Wagner Meira Jr.

Depart. of Computer Science
Federal University of Minas
Gerais
meira@dcc.ufmg.br

Ismael S. Silva

Depart. of Computer Science
Federal University of Minas
Gerais
ismael.silva@dcc.ufmg.br

Raquel O. Prates

Depart. of Computer Science
Federal University of Minas
Gerais
rprates@dcc.ufmg.br

Mohammed J. Zaki

Depart. of Computer Science
Rensselaer Polytechnic
Institute
zaki@cs.rpi.edu

Adriano Veloso

Depart. of Computer Science
Federal University of Minas
Gerais
adrianov@dcc.ufmg.br

Abstract

In this paper we describe the preliminary results and future directions of a research in progress, which aims at assessing the hashtag effectiveness as a resource for sentiment analysis expressed on Twitter. The results so far support our hypothesis that hashtags may facilitate the detection and automatic tracking of online population sentiment about different events.

Author Keywords

Twitter; Hashtag; User Sentiment Analysis; Social Phenomena

ACM Classification Keywords

H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous;

Introduction

Currently, determining people's sentiment about distinct subjects has become a strategic problem with applications in several scenarios, such as product marketing [7], election campaigns [4], and epidemics surveillance [1, 6]. A promising source of information for such analysis is social media (e.g., Twitter). The Twitter interaction model induces users to share and express continuously their opinions and sentiments, which are propagated to their followers. Further, users may post new messages

called tweets or share messages posted by someone else through retweets. In both cases, the content is bounded to 140 characters. However, determining the individual sentiment of each message coming in a large flow of data is laborious, error-prone, subject to ambiguity, and raises several research and technological challenges [9].

One strategy to address these challenges is to exploit hashtags contained in tweets. Hashtags are metadata expressed as a set of contiguous words preceded by the character #, and were originally proposed for organizing and grouping related messages¹. The increasing usage of hashtags enabled their use not only to organize content, but also to retrieve, propagate, promote and publicize a given content, or even create new trends and social mobilization [2]. In the face of the potential of hashtags for information diffusion, some researchers have been analyzing how this kind of metadata is being used from various perspectives [2, 10].

Considering hashtags original proposal and their current usage, we may hypothesize that they are also being used to express, group, and propagate people's sentiment with regard to some topics and events (e.g., "I watched a movie #Ilovelt"), and may contribute to detect and track sentiment. Such trend is interesting because hashtags may be useful to reduce the impact of ambiguity and linguistic variations, which are inherent to Twitter messages.

Motivation and goals: In order to investigate this hypothesis and motivated by the potential benefits from better sentiment analysis to several areas, we describe in this paper partial results and future directions of an ongoing research that aims at analyzing whether the

opinions expressed through Twitter hashtags (i.e., hashtags that represent sentiments and opinions, such as #Loveit) may support the analysis and tracking of sentiments with regard to specific topics (e.g., elections and diseases). Once we confirm the validity of hashtags for sake of sentiment analysis, we plan to use them as a source of information to automatically detect and track sentiment about specific events and topics that are discussed in large data flows in Twitter.

In face of our overall goal, we first address the following research question: Can we detect and track the online population sentiment from Twitter hashtags? The first step to answer this research question was to assess and characterize whether and how users employ hashtags to express their sentiments, how they propagate through the network, and to what extent they match actual sentiments. We perform this assessment based on a Twitter dataset related to the 2010 Brazilian presidential elections. The preliminary results support our hypothesis that hashtags may facilitate the automatic detection and tracking of online population sentiment about different events.

Research contributions: Our CHI-related contributions in terms of both preliminary results and overall goals are towards a better understanding of the use of hashtags that express sentiments in Twitter. Based on this understanding, new services, such as hashtag recommendation systems, may be designed in order to improve user experience, as suggested by Willinger et al. 2010 [11].

Furthermore, once we characterize the effectiveness of hashtags for sentiment analysis, this kind of metadata may be used for automatically detecting and tracking online population opinion. We emphasize that there is an

¹<http://support.twitter.com/articles/49309-what-are-hashtags-symbols>

Sentiment Polarity	Hashtag Sample (in Portuguese)	Hashtag Sample (translated to English)
Positive	#Brasil13	#Brazil13
	#dia31vote13	#Day31vote13
	#DiaDaDilma	#DilmaDay
Negative	#DilmaNão	#DilmaNot
	#votoNulo	#IVoteNull
	#OdeioDilma	#IHateDilma
Ambiguous	#Dilma13	#Dilma13
	#Marina43	#Marina43
	#Serra45	#Serra45
Neutral	#Eleições2010	#Elections2010
	#DebateGlobo	#DebateGlobo
	#Brasil	#Brazil

Figure 1: Example of Classified Hashtags

increasing demand for this type of application not only from actual users, but also from institutions that may use this information strategically, such as for disease surveillance.

Related Work

Some previous works focused on characterizing how information propagates in social media and its impact for various applications. In Jansen et al., 2009 [7], the authors analyzed the use of micro-blogging as an online mechanism for word-of-mouth marketing and discuss the implications for sake of a marketing strategy [7]. In Diakopoulos and Shamma, 2010 [4], the authors developed an analytical methodology and visual representations to better understand the temporal dynamics of sentiments expressed in social media.

Diakopoulos and Naaman et al., 2011 [3] described an ongoing research that aims at understanding and providing information that may help the moderation process of comments of online news [3]. Finally, in Miller et al., 2011 [8], the authors analyzed how opinions flow through connected posts. The study showed that the sentiment spreads in a cascade, indicating that people are influenced by the messages from their contacts.

Romero et al., 2011 [10] focus on verifying the differences among the various types of information with regard to spread pattern in social networks and found that the propagation patterns of hashtags and related topics are very similar. In Cunha et al., 2011 [2] a linguistic-inspired study is described to analyze how hashtags are created, used and disseminated. The research goal was to verify the hashtags' impact in the linguistic innovations. The results indicated that hashtags may effectively serve as models for characterizing the propagation of linguistic

forms [2].

Although these are related works, to the best of our knowledge, we did not find any works that focus on analyzing, characterizing and exploiting hashtag's potential for automatic detection and tracking of online population's sentiment about different events.

Data Collection

We started our analysis from tweets related to Dilma Rousseff in the Brazilian presidential election in 2010. This event was a major one in Brazil and reached high popularity on Twitter in 2010 ². Dilma Rousseff used this micro-blog as a main information source to her voters, achieved more than 500,000 followers and was the second most cited person on Twitter in 2010.

The dataset was collected using the Twitter Streaming API. Our dataset comprises 10,173,382 messages about the elections, from which 6,989,346 (69%) mentioned Dilma Rousseff and 1,760,303 (25%) contained some hashtag that could express sentiment or not. Furthermore, we identified 84,748 distinct hashtags, since the same hashtag might have been used more than once.

In our initial analysis, we focused on each hashtag without taking into consideration the tweet content to which it was assigned, in order to assess the expressiveness of the hashtag. In practice, two authors of this study classified manually the 10,000 most frequent hashtags. The hashtags were classified as positive or negative, when they expressed an explicit sentiment. For instance, the #DilmaNot hashtag indicates that users do not want Dilma as President, so it was classified as negative. Furthermore, hashtags could also be classified as

²<http://yearinreview.twitter.com/2010/trends/>




Figure 2: Frequency of positive and negative hashtags per day for Dilma Rouseff




Figure 3: Positivity of tweets mentioning Dilma across time

ambiguous - when the related tweet could be associated to a sentiment, but the hashtag did not make it clear, or neutral - when the hashtag was not associated to any sentiments. Figure 1 shows some samples of hashtags classified in each sentiment polarity.

Once we concluded the manual classification, we observed the existence of hashtags with the same content, however, written in different ways (e.g., uppercase, lowercase, with or without accent and underscore - #DilmaNot and #dilma_not). In such cases, the hashtags were unified and their frequencies were added.

At the end of this step, we verified that, from the 10,000 hashtags classified, 2,912 (29%) expressed some sentiment, positive (879 - 30%), negative (1,560 - 54%) or ambiguous (473 - 16%). It is noteworthy that in this stage of the study we analyzed only the hashtags that expressed an explicit sentiment (i.e., positive or negative). In the next section we describe the analysis performed and the preliminary results.

Analysis and Initial Observations

As mentioned, the first step of our work consisted of characterizing whether and how Twitter users are using hashtags to express, share and propagate their opinions about specific subjects. Towards that end, we analyzed (1) whether the opinion associated with hashtags matches the actual population sentiment across time and (2) the popularity and propagation pattern of the Twitter hashtags.

We started by performing an analysis of the use frequency of negative and positive hashtags at a daily basis for Dilma Rouseff, then candidate for the Brazilian presidency, as shown in Figure 2. We verified that, although some hashtags were used to express sentiments

in the beginning of the year, their use was really effective at a daily basis after the beginning of the campaign in standard media (i.e., radio and TV) on August 17th. The graph also shows the use of hashtags to express sentiments stabilizes again on November. This behavior is explained by the end of the elections, when Dilma won.

In order to better understand the results, the graph in Figure 3 presents the hashtag positivity for messages mentioning Dilma from August to November, 2010. Positivity is the ratio between the number of messages containing positive hashtags and the number of messages containing both positive and negative hashtags.

Through both graphs from Figures 2 and 3, we may verify that Dilma's positivity was low, close to the first election round (10/03/2010), matching the fact that she did not get enough votes to achieve absolute majority and to win in the first round, although she was the most voted candidate (47%). On the other hand, we may observe an increase in the positivity after October 16th, when the candidate Marina Silva (3rd in the first round) released her supporters to vote as they wished and a significant portion of them started supporting Dilma. Finally, one week before the second round, we observe again a decrease in the positivity, which was associated with polemical debates about topics such as abortion, religion, and same sex marriage, among others. Later, there were media campaigns to clarify her opinions and her positivity increased again, reaching an 80% overall maximum, and she won the elections on 10/31. In summary, the positivity behavior across time endorses the hypothesis that the hashtags sentiments match actual facts and the overall population sentiment.

We then performed the second step of our analysis, which analyzed the hashtags popularity and their propagation




Figure 4: Frequency ranking of hashtags in tweets that mentioned Dilma - A log-log plot showing volume of tweets in which the hashtag was used vs. its position in a popularity ranking

% of hashtags that appear in at least j -tweets				
Sentiment	$j = 100,000$	$j = 10,000$	$j = 1,000$	$j = 100$
Negative	0.06%	0.45%	1.60%	11.28%
Positive	0.00%	1.25%	4.66%	18.09%

Distribution of most popular hashtags that express sentiment

% of hashtags that appear among 6 and in at most i -tweets				
Sentiment	$i = 7$	$i = 8$	$i = 9$	$i = 10$
Negative	9.81%	16.03%	21.73%	28.21%
Positive	7.28%	10.81%	15.47%	21.62%

Distribution of less common hashtags that express sentiment

Figure 5: Hashtag probability distribution

process in Twitter. We want to verify whether the information propagation in social networks follows a cascade model. The cascade model for social networks, discussed in [5], states that people make their choices, consciously or not, based on someone else's sentiments and choices. The graph in Figure 4 presents the popularity ranking of each hashtag that expresses a sentiment, where the popularity is the number of hashtag occurrences in messages.

We verified that the frequency distribution of both positive and negative hashtags is very skewed, so that a small number of hashtags are very popular, which may be seen as "the rich get richer" [5], where few hashtags express the sentiment of most messages while the majority of the hashtags are used in just a few tweets. It is important to emphasize that a similar result was found when we consider linguistic variation [2].

We may better understand these findings by checking Figure 5, which presents the probability distribution of the most and the least popular hashtags, both for positive and negative sentiments. We observed that less than 10% of the hashtags, regardless of the sentiment, occur in more than 1,000 tweets.

These results also support the cascade propagation of sentiment-related hashtags, since a small group is employed in the majority of the messages, regardless of being positive or negative. This indicates that many people choose the same hashtags to express and propagate their sentiment. That is, consciously or not, they choose a hashtag to express their sentiment that was proposed or chosen by someone else to express theirs, as suggested in [5].

Since Twitter allows sharing ideas through both new

tweets or retweets, we also investigated how the sentiment-related hashtags propagated in the sharing process. It is important to emphasize that, in the first case (i.e., new tweets), users express their sentiment, or at least add an observation to someone else's opinion, while in the second case (i.e., retweet) they just endorse a message. We found that both strategies present the same popularity when verifying all the messages, but the popularity varies significantly across hashtags.

Finally, considering that hashtags are used in a significant number of new tweets and its increasing usage is a mechanism for both defining categories and spreading ideas [2], we performed a qualitative analysis of how these hashtags are used in composing a tweet. In particular, we verified whether the hashtag was used just as a tag for clustering and reinforcing the message sentiment or to assign sentiment to the message, which does not carry it without the hashtag.

We selected randomly 500 tweets that were evaluated by two of the authors of this paper. During the evaluation we verified whether: (1) the tweet sentiment matched hashtag sentiment and (2) just the message content was enough to determine the message sentiment. We found that just in 3% of the messages the hashtags contradict the message content. We then evaluated the remaining 97% regarding the contribution of hashtags and found that in 31% of the messages the hashtag was necessary for defining the message sentiment. Despite the significative findings, we believe that we need to deepen the analysis, which we intend to do as next steps of this research.

Final Remarks

Although preliminary, our results already describe how Twitter users have used hashtags to express, classify, and

propagate their opinions and sentiments. Furthermore, they support our hypothesis that hashtags do carry sentiments and the use of such information to portray the sentiment of a population about a topic is promising.

As next steps, we envisage activities such as characterizing the usage of hashtags for expressing sentiments related to other topics (e.g., disease outbreaks, product evaluation and sport events) as well as assessing the accuracy of hashtags to describe sentiments.

We employed these different perspectives to verify usage and propagation patterns of hashtags that may support the measurement and monitoring of online population opinion with respect to various subjects. Once we confirm that hashtags may be used for sentiment analysis, we may start using them for applications that aim to detect and track sentiments on specific topics and events that are discussed in large flows of data from social media.

Acknowledgments

We thank INCT-Web(INWeb),FAPEMIG, CNPq, CAPES, RPI and UOL, for the support to this research.

References

- [1] C. Chew and G. Eysenbach. Pandemics in the Age of Twitter: Content Analysis of Tweets during the 2009 H1N1 Outbreak. *PLoS ONE*, 5(11):e14118+, 2010.
- [2] E. Cunha, G. Magno, G. Comarela, V. Almeida, M. A. Gonçalves, and F. Benevenuto. Analyzing the dynamic evolution of hashtags on twitter: a language-based approach. In *Proc. of the Workshop on LSM*, pages 58–65, 2011.
- [3] N. Diakopoulos and M. Naaman. Topicality, time, and sentiment in online news comments. In *Proc. EA CHI*, pages 1405–1410, 2011.
- [4] N. A. Diakopoulos and D. A. Shamma. Characterizing debate performance via aggregated twitter sentiment. In *Proc. CHI*, pages 1195–1198, 2010.
- [5] D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, 2010.
- [6] J. Gomide, A. Veloso, W. M. Jr., V. Almeida, F. Benevenuto, F. Ferraz, and M. Teixeira. Dengue surveillance based on a computational model of spatio-temporal locality of twitter. In *ACM WebSci*, 2011.
- [7] B. J. Jansen, M. Zhang, K. Sobel, and A. Chowdury. Micro-blogging as online word of mouth branding. In *EA CHI*, pages 3859–3864, 2009.
- [8] M. Miller, C. Sathi, D. Wiesensthal, J. Leskovec, and C. Potts. Sentiment Flow Through Hyperlink Networks. In *Proc. AAAI CWSM*, 2011.
- [9] B. Pang and L. Lee. Opinion mining and sentiment analysis. *Found. Trends Inf. Retr.*, 2:1–135, 2008.
- [10] D. M. Romero, B. Meeder, and J. Kleinberg. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proc. WWW*, pages 695–704, 2011.
- [11] W. Willinger, R. Rejaie, M. Torkjazi, M. Valafar, and M. Maggioni. Research on online social networks: time to face the real challenges. *SIGMETRICS Perform. Eval. Rev.*, pages 49–54, 2010.