Classifying Trending Topics: A Typology of Conversation Triggers on Twitter

Arkaitz Zubiaga
NLP & IR Group @ UNED
C/Juan del Rosal, 16
Madrid, Spain
azubiaga@lsi.uned.es

Damiano Spina
NLP & IR Group @ UNED
C/Juan del Rosal, 16
Madrid, Spain
damiano@lsi.uned.es

Víctor Fresno
NLP & IR Group @ UNED
C/Juan del Rosal, 16
Madrid, Spain
vfresno@lsi.uned.es

Raquel Martínez
NLP & IR Group @ UNED
C/Juan del Rosal, 16
Madrid, Spain
raquel@lsi.uned.es

ABSTRACT

Twitter summarizes the great deal of messages posted by users in the form of trending topics that reflect the top conversations being discussed at a given moment. These trending topics tend to be connected to current affairs. Different happenings can give rise to the emergence of these trending topics. For instance, a sports event broadcasted on TV, or a viral meme introduced by a community of users. Detecting the type of origin can facilitate information filtering, enhance real-time data processing, and improve user experience. In this paper, we introduce a typology to categorize the triggers that leverage trending topics: news, current events, memes, and commemoratives. We define a set of straightforward language-independent features that rely on the social spread of the trends to discriminate among those types of trending topics. Our method provides an efficient way to immediately and accurately categorize trending topics without need of external data, outperforming a content-based approach.

Categories and Subject Descriptors

H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval; H.1.2 [**Models and Principles**]: User/Machine Systems—*Human information processing*

General Terms

Human Factors, Experimentation, Measurement

Keywords

twitter, trending topics, social media, real-time, classification

CIKM'11, October 24–28, 2011, Glasgow, Scotland, UK.

1. INTRODUCTION

Twitter[1] is a microblogging site where users share short messages so-called tweets. These tweets tend to spread to a large number of users in very little time. Users on Twitter not only tweet about their personal issues or nearby events, but also about more general topics or news [8]. Due to the large amounts and diversity of real-time information contained on the site, Twitter lists a freshly updated set of trending topics. Trending topics comprise the top terms being discussed currently on Twitter. This list of top terms, which is updated in real-time, provides a reflection of the current main interests of the community, i.e., the most-discussed conversations right at the moment. It has become an appealing feature for Twitter users, real-time application developers, and social media researchers, thank to the ability to detect trending topics in the earliest stage.

The interests of the large community of users on Twitter encompasses a vast number of topics and types of messages. Hence, there is not a single reason that makes a topic become a trend. We believe that detecting the type of a trending topic helps better understand the origin that made it trend. This detection enables to better harness real-time data, since each type of trending topic may be useful for a different purpose. Discriminating trending topics by type of trigger in real-time could be useful to filter out unwanted ones in information filtering systems, or to enhance real-time data processing by treating them separately. We found little work dealing with the nature and characteristics of trending topics and, as far as we know, no research has been conducted dealing with the immediate characterization and detection of types of trending topics.

In this work, we introduce a typology of 4 types of trending topic triggers: news, current events, memes, and commemoratives. We propose a set of 15 straightforward features that characterize the social spread of a trending topic. These features are independent of the languages in which tweets are written. Using a collection of trending topics categorized according to the above typology, we perform classification experiments using Support Vector Machines to find the extent to which these features can help automatically categorize trending topics. Our method presents the ad-

[1]http://twitter.com/

vantages that (i) it provides an accurate classification outperforming the baseline approach relying on the content of tweets, (ii) it only requires a small set of features, which do not increase as the collection grows, (iii) it does not require using any external data, and (iv) its low computational cost enables to immediately predict the type of a trending topic as it emerges and appears on Twitter's homepage.

Next, in Section 2 we describe Twitter's trending topics. In Section 3 we summarize the related work. Then, in Section 4 we introduce a typology to organize trending topics by type of triggers, and detail the dataset used in the experiments. We propose a set of features, describe the trend classification experiments, and discuss the results in Section 5. Finally, we conclude the work in Section 6.

2. TRENDING TOPICS

One of the main features on the homepage of Twitter shows a list of top terms so-called trending topics at all times. These terms reflect the topics that are being discussed most in the latest minutes on the site's stream of tweets. In order to avoid topics that are always popular, Twitter focuses on topics that are being discussed much more than usual, i.e., topics that recently suffered an increase of use, so that it trended for some reason. Trending topics have generated big interest not only for the users themselves but also for information seekers such as journalists, real-time application developers, and social media researchers. Being able to know the top conversations being discussed at a given time helps keep updated about current affairs, and discover the main concerns of the community. Twitter defines trending topics as *"topics that are immediately popular, rather than topics that have been popular for a while or on a daily basis"*[2]. A trending topic is made up by the topic itself –i.e., the term that became a trend–, and a stream of tweets containing that topic.

3. RELATED WORK

So far, most of the work on Twitter has focused on analyzing the microblogging phenomenon [5, 7], modeling the information diffusion on the social network [13, 14] and analyzing the content of tweets [12, 10, 4].

Regarding the classification of single tweets, Sriram et al. [12] define a typology of five generic classes of tweets (news, events, opinions, deals, and private messages) in order to improve information filtering. The authors represent tweets using a small set of language-dependent features to classify tweets written in English. The use of these features outperforms the bag-of-words approach in the classification of tweets according to the typology. We believe that this typology, specifically defined for single tweets, does not fit the nature of triggers that leverage trending topics. Furthermore, this typology does not intend to classify by type of trigger. For instance, it is not intuitive to include a trending topic produced by a memorial day in this typology.

Little work has been done analyzing the properties of trending topics on Twitter. Most of them focus on event and topic detection [1, 2, 10]. Asur et al. [1] explore the longevity of trending topics on Twitter, and analyze the role of users in the emergence of trends. They found that (i) a few users are *trend-setters*, i.e, early contributors in the emergence of a trend, and (ii) a larger group of users are *propagators*, who help spread the topic. Cheong and Lee [3] and Kwak et al. [7] analyze the evolution of trending topics over the time, and perform a qualitative study of social features that characterize trending topics. They rely on the whole lifetime of trending topics to perform this analysis. Different from these, our work aims at immediately detecting the type of a trending topic as it emerges, so that we do not consider the whole lifetime of a trend.

Different from the above research, we present a typology to organize trending topics by trigger, and we a propose a set of language-independent features, and introduce a method that processes trending topics as soon as they emerge.

[2]http://support.twitter.com/articles/101125-about-trending-topics

4. TYPOLOGY OF TRENDING TOPICS

Next, we introduce the typology we use in this work to organize trending topics. After describing the typology, we detail the process of generation of a dataset, made up by trending topics and organized according to the typology.

4.1 Definition of a Typology of Trending Topics

Different happenings, either in the world, on TV, or on the Internet, can motivate users to discuss on a topic. After tracking trending topics for months, and acquiring insight on their emergence, we set forth a typology composed by the following 4 triggers:

- **News:** users sharing a breaking news promptly give rise to a trending topic. It has been shown that sometimes news break on Twitter before other online news media [7]. Thus, the detection of trending topics produced from news can be of utmost interest to discover breaking news. The early detection of news can also help feed news curation services, and notify search engines to update indexes related to the topic.

- **Current event:** users commenting on an event that is currently taking place leverage a trending topic in many cases. Users tend to live tweet about an event they are following (e.g., a soccer game or a TV show) or attending (e.g., a festival or a conference). The detection of these trending topics can help discover events that are taking place, or discover the events in which users are interested most.

- **Meme:** memes can also become a trending topic on Twitter. Memes are ideas that propagate through a social network. Usually, it is an idea that a community or a single (usually influential) user decides to launch in order to make it grow rapidly, and make it viral. The idea behind a meme may vary; it may be a community's will for something to happen, or a funny idea that someone proposed. The detection of these trending topics is useful to study the virality of ideas on the Internet, and for users looking for funny stuff.

- **Commemorative:** commemorative tweets can also become a trending topic. We consider it a commemorative topic when users congratulate a celebrity for his birthday, share about the anniversary of an event, or it is a memorial day. These trending topics can sometimes even be predicted, since upcoming celebrations and commemorations can previously be known. However, detecting them can be useful to analyze the way they come out, how they spread, and what opinions or comments users attach.

4.2 Dataset

Twitter selects 10 trending topics that are being discussed most at the moment. Using Twitter's top trends and search API methods, we monitored the trending topics shown on the site from March 1st to 7th. The list of top 10 trending topics was requested every 30 seconds. Thus, the process guarantees the detection of a trending topic almost as soon as it appears on the site, with a delay of 30 seconds in the worst scenario. As soon as a new topic appeared in the list of trends, another process queried the search API for the latest tweets containing the topic as a query term. Following this process, we collected a total of 1,036 unique trending topics. These trends include a total of 567,452 tweets from 348,757 different users. Accordingly, each of the 1,036 trending topics in the dataset contains an average of about 548 associated tweets. All these tweets are written in 28 different languages, with a majority of 295,082 tweets written in English, 76,628 in Spanish, 67,673 in Portuguese, 31,685 in Dutch, and 22,863 in Indonesian. Moreover, there is not a trending topic with just one language in the dataset.

All these topics were manually categorized within the typology we defined: news, current events, memes, or commemoratives. Each trending topic was included into just one type. To provide the annotations, the stream of tweets corresponding to a trending topic was read carefully. After understanding the trigger that caused the conversation to become a trending topic, the annotation was provided. Due to the large number of languages contained in the collection, machine translators were used for unknown languages. The annotation produced an organization of the 1,036 trending topics in groups of 616 current events, 251 memes, 142 news, and 27 commemoratives.

5. CLASSIFYING TRENDING TOPICS

Next, we present the features to characterize trending topics, and the classification experiments.

5.1 Features of Trending Topics

As an approach to discovering the type of a trending topic, we propose 15 social features that consider the way it spreads. Furthermore, since we want to categorize a trending topic as soon as it appears trending on the site, the features must be straightforward, easy to get, and cheap to compute while the system performs accurately. This would ensure the immediacy of computation, and the ability to predict the type of a trending topic on the fly, as soon as it appears on the Twitter's list. Moreover, these features are independent of the language used in tweets, and do not depend on the vocabulary utilized by users.

On one hand, we use average number of occurrences of features in the tweets corresponding to a trend. Each average computed as the arithmetic mean is the result of dividing the number of occurrences of the corresponding feature in the whole trend by the total number of tweets gathered for the trending topic. Note that the number of tweets in a topic that just trended is relatively small (average of 548 tweets). We propose 10 different features that rely on average values: (1) Level of retweets (the number of retweeting users to reach the current state of a tweet), (2) Ratio of retweets out of all tweets, (3) Hashtags, (4) Length of tweets, (5) Exclamations, (6) Questions, (7) Links, (8) Repetition of the trending topic in tweets, (9) Replies, and (10) Spread velocity (tweets per second).

On the other hand, we compute the diversity of feature values all across the tweets in a trending topic. The diversity calculates the variation of the feature throughout the trending topic. The higher is the diversity value, the more different is the feature from tweet to tweet within a trending topic. To compute the diversity, we use the Shannon's diversity index [11]. We propose 5 features that rely on diversity: (1) Contributing user diversity, (2) Retweeted user diversity, (3) Hashtag diversity, (4) Language diversity, and (5) Vocabulary diversity.

5.2 Classification of Trending Topics

5.2.1 Experimental Setup

We perform classification experiments according to the established typology using Support Vector Machines (SVM) [6] as a state-of-the-art classification algorithm. Being a multiclass task, we rely on a *one-against-all* binary combination method [9] using *svm-light*[3]. We set the SVM to run with a linear kernel and the default parameters.

We use two different representations in the classification process: the Twitter features described in Section 5.1, and a bag-of-words of the textual content of tweets as a baseline. We create one vector per trending topic for each of the representations. We use a training set with 600 trending topics, with 436 in the test set. We perform 10 random selections of the training set, and show the average results. The two representations we use are the following:

- **Twitter features:** we use vectors with 15 dimensions, corresponding to the features introduced above. The main advantages of this representation approach are the straightforwardness of the features, and that the number of features does not depend on the number of trends we have to represent, but it remains unchanged.
- **Textual content:** we rely on the bag-of-words approach. We represent each trend with the TF values of each term appearing in the stream after removing stopwords. This approach is computationally much more expensive than the approach based on Twitter features, and presents the problem that the number of dimensions utilized to represent the trends increases as the collection grows. In this case, representing the 1,036 trending topics generated vectors with 512,943 dimensions.

5.2.2 Evaluation

	Accuracy	Kappa
Bag-of-words	.752	.530
Twitter features	.784	.604

Table 1: Accuracy and Cohen's Kappa statistic of the trend classification by type of representation

Table 1 shows the accuracy values of the classification of trending topics. The accuracy measures the percent of correct guesses among all the predictions. The representation

[3]http://svmlight.joachims.org

	Class Precision			
	N	**CE**	**M**	**C**
Bag-of-words	.673	.783	.636	.548
Twitter features	.634	.829	.731	.132

Table 2: Precision by class (N: News; CE: Current Events; M: Memes; C: Commemoratives) of the trend classification by type of representation

based on the proposed Twitter features achieves a superior accuracy than the bag-of-words baseline. This superiority gap is of 3.2% in favor of the representation based on Twitter features. Furthermore, the winning approach presents the advantages that the features are straightforward and easy to compute, and that it only requires 15 features instead of the thousands required by the content-based representation. Cohen's kappa values in Table 1 show the impact of randomness on the agreement between predicted and observed classifications. Note that the higher is the value, the lower is the effect of randomness on the accuracy improvement. It can be observed that the kappa values follow the same rank as the accuracy values, so that it indicates that the accuracy improvement is not a consequence of random guesses.

When analyzing in more depth the precision by class (see Table 2), it can be seen that Twitter features perform better in the detection of current events, and especially memes, where the gap is bigger than 9%. Using the bag-of-words improves for news, and especially for commemoratives, where Twitter features get low precision. However, this could be due to the small number of commemoratives in the collection[4], and it may perform better with more trends of this type, probably providing a better representation of the class in the training stage.

6. CONCLUSIONS

In this work, we have explored the types of triggers that leverage conversations on Twitter. We have introduced a typology to organize Twitter's trending topics by the type of happening that caused them. This typology includes the following 4 types of trending topics: *news*, *current events*, *memes*, and *commemoratives*.

We aimed at characterizing trending topics so that we were able to organize them by type. To this end, we have set out 15 straightforward features to represent trending topics. These features are independent of the language used in the tweets, rely on the social diffusion of the trending topic, and only require tweets sent before the topic trended, without need of external data. The proposed method provides an immediate way to accurately organize trending topics. Unlike for the content-based representation, the number of features remains unchanged, and does not increase as the collection grows. Using SVM classifiers upon these features to discriminate types of trending topics showed that the proposed features provide more accurate classification results than the use of textual content of tweets.

To the best of our knowledge, this is the first research work introducing a typology of trending topics, and providing a method to immediately classify trending topics as soon as they appear on the homepage of Twitter. This study is relevant not only for researchers studying trends in social media, but also for developers who enhance their systems with real-time data from social networks. Furthermore, it paves the way to researchers interested in exploring causes that leverage trends on social media.

[4]Note that there are only 27 commemorative trends in the collection.

7. ACKNOWLEDGMENTS

This work has been part-funded by the Education Council of the Regional Government of Madrid, MA2VICMR (S-2009/TIC-1542), the Innovation project Holopedia (TIN2010-21128-C02-01), and the Spanish Ministry of Education for a doctoral grant (AP2009-0507).

8. REFERENCES

[1] S. Asur, B. A. Huberman, G. Szabó, and C. Wang. Trends in social media : Persistence and decay. *CoRR*, abs/1102.1402, 2011.

[2] M. Cataldi, L. Di Caro, and C. Schifanella. Emerging topic detection on twitter based on temporal and social terms evaluation. In *Proc. of MDMKDD '10*, pages 4:1–4:10. ACM, 2010.

[3] M. Cheong and V. Lee. Integrating web-based intelligence retrieval and decision-making from the twitter trends knowledge base. In *Proc. of SWSM '09*, pages 1–8. ACM, 2009.

[4] J. Hurlock and M. L. Wilson. Searching twitter: Separating the tweet from the chaff. In *Proc. of ICWSM '11*. AAAI, 2011.

[5] A. Java, X. Song, T. Finin, and B. Tseng. Why we twitter: understanding microblogging usage and communities. In *Proc. of WebKDD/SNA-KDD '07*, pages 56–65. ACM, 2007.

[6] T. Joachims. Text categorization with support vector machines: learning with many relevant features. In *Proc. of ECML '98*, pages 137–142, 1998.

[7] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In *Proc. of WWW '10*, pages 591–600. ACM, 2010.

[8] E. Mishaud. Twitter: Expressions of the whole self. Master's thesis, Department of Media and Communications, University of London, 2007.

[9] R. Rifkin and A. Klautau. In defense of one-vs-all classification. *Journal of Machine Learning Research*, 5:101–141, December 2004.

[10] T. Sakaki, M. Okazaki, and Y. Matsuo. Earthquake shakes twitter users: real-time event detection by social sensors. In *Proc. of WWW '10*, pages 851–860. ACM, 2010.

[11] C. E. Shannon. A mathematical theory of communication. *SIGMOBILE Mob. Comput. Commun. Rev.*, 5:3–55, January 2001.

[12] B. Sriram, D. Fuhry, E. Demir, H. Ferhatosmanoglu, and M. Demirbas. Short text classification in twitter to improve information filtering. In *Proc. of SIGIR '10*, pages 841–842. ACM, 2010.

[13] S. Wu, J. M. Hofman, W. A. Mason, and D. J. Watts. Who says what to whom on twitter. In *Proc. of WWW '11*, pages 705–714. ACM, 2011.

[14] J. Yang and J. Leskovec. Modeling information diffusion in implicit networks. *Proc. of ICDM '10*, 0:599–608, 2010.