#tag: Meme or Event?

Dimitrios Kotsakos
University of Athens, Greece
Email: dimkots@di.uoa.gr

Panos Sakkos
University of Athens, Greece
Email: psakkos@di.uoa.gr

Ioannis Katakis
University of Athens, Greece
Email: katak@di.uoa.gr

Dimitrios Gunopulos
University of Athens, Greece
Email: dg@di.uoa.gr

***Abstract*—Users in social networks use hashtags for various reasons, some of them being serving search purposes, gaining attention or popularity or starting new conversation - thus, creating viral memes. In this paper we address the problem of classifying these hashtags in different categories, based on whether they represent a real life event or a social network generated meme. We compute a set of language-agnostic features to aid the classification of hashtags into events and memes and we provide an extensive study of the behavior that characterizes memes and events. We focus on Twitter social network, we apply our methods on a big dataset and reveal interesting characteristics of the two classes of hashtags.**

I. Introduction

On-line social networks analysis recently attracted attention from many scientific fields. In many cases, research on social data is interdisciplinary. This constantly raising interest is certainly expected, since social network data are easy to access and reflect various aspects of human behaviour and community dynamics. Probably the most well studied social network, is the micro-blogging platform Twitter.

New data mining tasks have recently appeared in micro-blog environments presenting interesting research challenges as well as commercial value. Sentiment Analysis [1], Event Recognition [2], Trend Identification [3], Community Recognition [4], Influence Propagation [5] are just a few characteristic examples.*Tagging* thrives in Web sites with user-submitted content where tags are voluntarily assigned for information retrieval purposes, since tag-based search functionality is available. Tags are assigned to many different types of information sources such as images, videos and music. Twitter is a tag-rich service since users annotate Tweets by inserting keywords that are marked with the hash (#) character. These keywords are known as *hashtags* (see Figure 1). Hashtags are considered very important keywords since they add valuable meta-knowledge to text that is limited to 140 characters. In order to track certain events and to annotate them properly users agree to hashtag Tweets with a predefined keyword (e.g. `#asonam14`). Many micro-blog analysis tasks are exploiting tagging behaviour. Hence, hashtag quality plays an important role for information organization.

In social media, users use hashtags not only to annotate specific events but also to promote ideas or discussions known as *internet memes*. Many times Memes arise when a group of Celebrity fans try to promote a discussion topic related to their pop idol (Figure 2). Other types of Memes include internet hoaxes or marketing material. Memes are not inherently detrimental. However, since their data volume is many times significant, they can obstruct other tasks like trend or event detection, where Memes are considered as noise.

Fig. 1. A Tweet that utilizes hashtags to annotate content

Fig. 2. A Tweet promoting a Meme

Our contributions can be summarized as follows: *a)* We provide a definition of *meme* and *event* and discriminate between them in social networks, *b)* we propose a set of language-agnostic features to aid the classification of hashtags into *Event* or *Meme* and *c)* we provide an extensive study of the behavior that characterizes *memes* and *events* and present an Gain-Ratio-based ranking of the proposed features in our setting.

II. Related Work

In this section, we discuss research efforts that study a) Memes, b) Trend and Event detection, and c) Hashtags.

a) Memes: Bauckhage [6] defines internet Memes as evolving content that rapidly gains popularity or notoriety on the Internet. Moreover, the author states that Memes are spread voluntarily rather than in a compulsory manner, which fact, although true, does not describe the full picture. Very often, Memes are produced by advertising or community campaigns, so they are expected to have different behavior to organically and not strategically created memes. The related bibliography lacks methods of recognizing these campaigns. The current paper offers an initial approach towards this direction. Leskovec et al. define memes as “short, distinctive phrases that travel relatively intact through on-line text” [7].

b) Trends and Events: In [8] the authors employ time series clustering in order to uncover temporal patterns in the popularity of content in social media and focus on the propagation of hashtags on Twitter. The authors (as in [7]) claim that mainstream media accounts (CNN, BBC, etc) produce content and push it to the other contributors, including twitter first consume-then produce accounts and professional bloggers.

c) Hashtag Analysis: Many approaches in social network mining research have been devoted to the analysis of hashtags in social networks. The hash symbol (‘#’) has been used to indicate the special meaning of a word and tag content in social networks like *Twitter, Instagram* or *Facebook*. Users use hashtags for search, annotations, or viral conversations often called *memes*. As opposed to traditional web search, queries in Twitter search that contain the hash symbol a significant portion of the total queries issued to the system [9]. Moreover, many Twitter queries reference words used in hashtags, but without the preceding '#' in the query. Since the amount of possible hashtags a user can use to either tag content or search for results is essentially unlimited, both these tasks would benefit if users were aware of tags used by other users [10]. Interestingly however, related work has not focused on the problem of distinguishing between *memes* and *events*.

III. Events vs Memes: Problem Definition

Memes vs Events. Memes and Events can be observed in a social stream of content s by identifying an excessive appearance of content related to this Meme or Event. The difference between an Event and a Meme is that an Event can be traced back in a news stream n of the same period (as in s) whereas a Meme only appears in s.

In both cases, Memes and Events, as with any other content, are annotated with hashtags. For example, Event hashtags could be: `#GermanyElections`, `#earthquake`, `#Oscars2014`, whereas for Meme hashtags could be the following: `#loveit`, `#insomnia`, `#20ReasonsIAmCute`, `#WantJustinInIreland`. As we can observe, there are not any structural characteristics that can aid in separating Events from Memes. In this work, we try to automatically build a model that distinguishes between the two.

Problem Definition: *Given a limited part of the social stream $s_T \subset s$ (training set), build a model that can assign a label (*`event, meme`*) to a hashtag h given a specific set of information for this hashtag.*

This set of information can be represented as a feature vector $\vec{h_x}$. The requested model is a function $f(\vec{h_x}) \rightarrow \{event, label\}$, where $\vec{h}_x = \{g_1(s_T, h_x), \dots, g_n(s_T, h_x)\}$ and g_i $(i = 1, \dots, n)$ is the function that calculates the value of feature i for hashtag h. We formulate this problem as a data classification problem by training traditional classifiers to learn the features that separate the two classes.

IV. Our Approach

A. Discussion

We propose a set of features that given a set of predefined classes and a manually labeled training set, can be used by a classifier with the goal of classifying the testing examples to the classes mentioned above. The classification task can include the classification of hashtags, topics or keywords. In this work, we focus on the hashtag classification, although our work can be applied in the classification of each of the above mentioned types.

We computed 15 different features, resulting in 15-dimensional vectors representing the topics (in this case hashtags) in the dataset. Some of the features are Twitter-specific (e.g. retweets or favorites), however they can be applied on all social networks that support sharing or promoting content (e.g. *Share* or *Like* in *Facebook*). For each hashtag h_i, we take its *hashtagLength* in characters into account. Communities aiming to promote a meme or a phrase for an advertising campaign often try to collapse a whole phrase into a single word in order to save characters. In this sense, often Memes are longer than Event-hashtags, because users try to embed in them as much information as possible. For example, not many English words are as long as `#WeWantOneDirectionInLondon`, which appeared in Twitter stream sometime in March, 2013. For each hashtag h_i we computed the following features:

1) Document features: The following features are computed over the set of all documents that contained hashtag h_i. We tried to capture the significance of rich content in tweets, e.g. links, pictures, or videos. We computed avarage values of the following features for each hashtag: *tokensPerTweet*, *hashTagsPerTweet*, *urlsPerTweet*, *mediasPerTweet*, *favoritesPerTweet*, and *retweetsPerTweet*.

2) Social features: With the social features we try to capture the importance of conversations in the social network. In our analysis we used *tweetsWithReplies*, which reflects the percentage of tweets with hashtag h_i that were replies to other tweets, and *mentionsPerTweet*, which is the avg. number of mentions to other users over all tweets with hashtag h_i.

3) Community features: Memes are expected to come from clusters of users, whereas events are expected to interest a broader user base. In this sense, we computed *uniqueUsersCount*, which represents the size of the community that was interested in the hashtag h_i. Moreover we used *statusesPerUser* as the avg. number of tweets from the set of unique users that posted a tweet with hashtag h_i. This feature captures the historical activity of the community that was active with respect to h_i. In order to capture some graph-related characteristics we used *userFollowersPerUser, userFriendsPerUser, listedCountPerUser*. These features capture the popularity and the social activity of the users that appeared to be interested in hashtag h_i. In order to have a measure of the credibility of the users interested in each hashtag, we utilize the *Verified* feature of Twitter, and we compute the avg. number of users that are verified by Twitter, as *avgVerifiedUsers*.

B. Dataset Description

We crawled Twitter using the Twitter Streaming API for the period between *February 16, 2014* and *April 6, 2014*. We collected tweets from the bounding box of United Kingdom and used the top-20 most popular hashtags per day, resulting in 27 million tweets written by $721,644$ users. Around 4 million tweets contained at least one of the 1.1 million unique hashtags. Figure 3 illustrates the distribution of the hashtags. It is apparent that most hashtags are used only once, which indicates that the users are not aware about which hashtags are used by other people in the network at the same moment. Our dataset contained hashtags about both *Memes* (e.g. `'georgesnapchatme'`, `'100happydays'`, etc.) and `Events` (e.g. `'brits2014'`, `'bbcqt'` tagging tweets about the BRIT Awards 2014 and the 'BBC Question Time' television program respectively, etc.). We asked 5 people to manually tag all hashtags into one of the two classes. In order to avoid bias towards any of the classes the annotators were not exposed to the feature vectors. Afterwards, we used majority voting in order to specify the class of each hashtag.

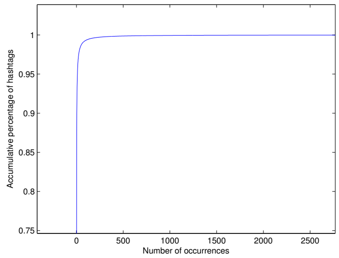

Fig. 3. Cumulative distribution of unique hashtags over number of occurrences

TABLE I. CONFUSION MATRIX OF THE FOUR CLASSIFIERS (M=MEME, E=EVENT)

	Prediction							
	Naive Bayes		**Random Forest**		**SVM**		**k-NN**	
True	M	E	M	E	M	E	M	E
M	0.64	0.36	0.91	0.09	0.73	0.27	0.87	0.13
E	0.08	0.92	0.11	0.89	0.13	0.87	0.14	0.86

We ended up with 1100 tagged vectors, among of which 558 where tagged as *Events* and 542 as *Memes*.

C. Experimental Study on Used Classifiers

We experimented with four traditional general purpose classifiers offered by the Weka tool [11]. Specifically, we chose the Naive Bayes, Random Forest, Support Vector Machines (SVM) and k-Nearest Neighbor classifiers (k-NN) [12]. It has been shown that Naive Bayes is effective in practice without the unrealistic independence assumption. The Random Forest classifier is effective in giving estimates of what variables are important in the classification, thus providing an importance ranking of the features [13]. The Support Vector Machine implementation we chose was the Sequential Minimal Optimization (SMO) algorithm [14], which trains a support vector machine with polynomial or RBF kernels.

Random Forest was able to reach an accuracy of 89.2%. The confusion matrices for all classifiers are illustrated in Table I. Figure 4 illustrates how the classifiers compare against each other in terms of accuracy as a function of training set size. Random Forest was the most accurate in all cases. Figure 4 compares the classifier accuracy with a 10-fold cross-validation scheme. Again, Random Forest outperforms the others reaching 89.66%.

D. Experimental Study on Feature Selection

In order to argue about which features are the most important for hashtag classification we ranked them in decreasing Gain Ratio with respect to the two classes II. We repeated the classification process, starting with the first feature and incrementally adding the rest one by one. Figure 4 depicts the results of this experiment. Here again, Random Forest outperforms all others for all subsets. Interestingly, when we used only *community* features, the Random Forest classifier was able to reach an accuracy of 70.8%, while when we used only *document* features the classifier reached 86% accuracy.

TABLE II. DECREASING GAIN RATIO FEATURE RANKING

Feature	**Gain Ratio**	**Feature**	**Gain Ratio**
1. tweetsPerUser	0.1617	8. userFriendsPerUser	0.0802
2. tweetsWithReplies	0.1432	9. mediasPerTweet	0.0778
3. userStatusesPerUser	0.1181	10. uniqueUsersCount	0.0574
4. tweetsWithUrl	0.0958	11. hashtagLength	0.0572
5. urlsPerTweet	0.091	12. hashTagsPerTweet	0.0527
6. tokensPerTweet	0.0822	13. avgVerifiedUsers	0.0461
7. mentionsPerTweet	0.0822	14. userFollowersPerUser	0.0355

Observations:

1) Tweets that contribute to propagation or promotion of *Memes* have significantly more videos, and photos (in general media) attached to them than tweets discussing real-life *Events*. Memes often are parts of campaigns or internet petitions and users try to enrich the content they generate so it ranks higher in search results, either for a specific hashtag or for a relevant topic (Figure 5).

2) Tweets that are relevant to *Memes* draw more conversations in the social network than tweets that report a real-life *Event*. As described above, the number of unique users who are interested in memes is relatively small and thus communities with similar meme-oriented interests are more easily formed. Such communities consist of people interested in celebrities, jokes, etc. (Figure 5).

3) Tweets discussing *Events* have on average slightly more tokens. This is normal, since these kind of tweets have a less arbitrary structure as they often include quotes or headlines in order to reproduce news reports, thus more words are needed to express something news-worthy (Figure 5).

4) Interestingly enough, Figure 6 reveals a rather odd observation. While tweets about breaking and significant events were expected to contain a relatively high number of URLs linking to external sites with the source of the information, this appears not to be true. In the Figure, there is a clear separation of the spaces covered by *Memes*-examples and *Events*-examples, showing that *Memes* are represented by tweets with fewer tokens - as described above - and more URLs, whereas *Events*-related tweets contain on average and on aggregate much fewer URLs and more tokens.

V. CONCLUSION AND FUTURE WORK

In this paper we defined the problem of distinguishing a popular topic of interest in a social network between network-generated topics of discussion, denoted as *Memes* and real-life events that triggered the interest of the social network users, denoted as *Events*. We provided a detailed study of the features that affect the classification, applying our experiments on the Twitter network using a dataset with 27 million tweets and 1.1 million unique hashtags. We used traditional classification methods, among of which the Random Forest classifier performed best, having beenable to reach an accuracy of 89% in its prediction on whether a topic is a *Meme* or an *Event*. Our study reveals interesting characteristics of the two classes of hashtags, some expected and some not. In our future work we plan to further investigate the features that characterize the behavior of popular topics and to create taxonomies of hashtags that facilitate recommendation or search tasks.

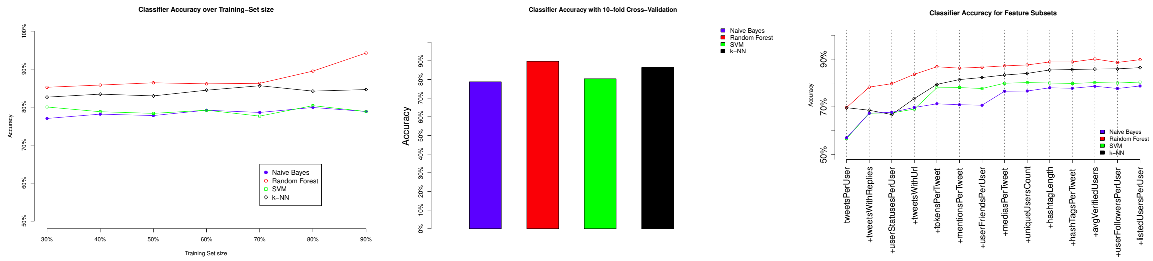

Fig. 4. Accuracy of the used classifiers over training set size, 10-fold cross-validation and incrementally adding features w.r.t to Gain Ratio ranking (Table II)

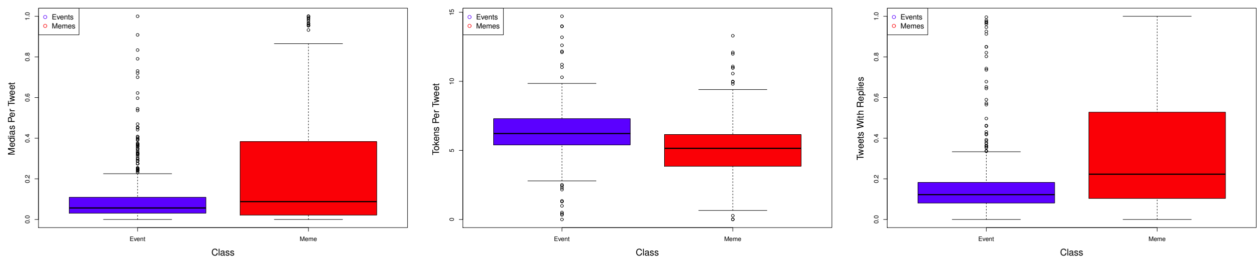

Fig. 5. Boxplots for the avg. number of hashtags, media entities, tokens and replies per relevant tweet against the two classes

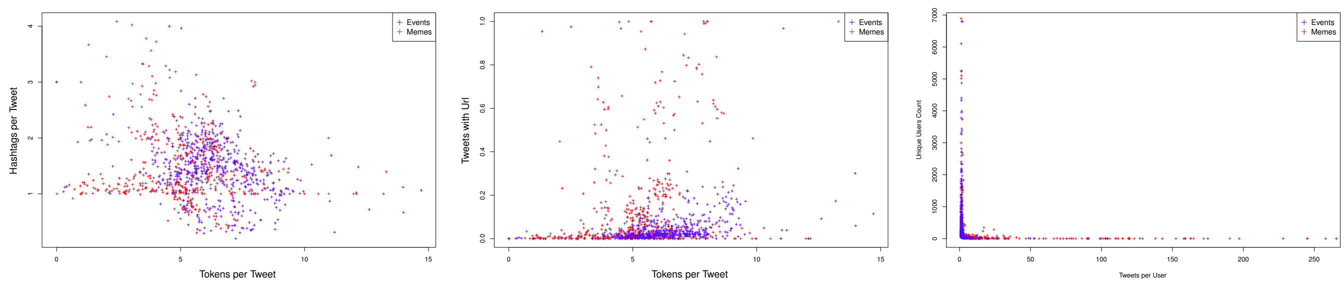

Fig. 6. Relationship between different pairs of features. *a)* Hashtag count and Tokens per tweet, *b)* URLs count and Tokens per tweet and *c)* Unique Users count and Tweets count per user

ACKNOWLEDGMENTS

This work has been co-financed by EU and Greek National funds through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) - Research Funding Programs: Heraclitus II fellowship, THALIS - GeomComp, THALIS - DISFER, ARISTEIA - MMD and the EU funded project INSIGHT.

REFERENCES

[1] E. Kouloumpis, T. Wilson, and J. Moore, "Twitter sentiment analysis: The good the bad and the omg!" *ICWSM*, vol. 11, pp. 538–541, 2011.

[2] G. Valkanas and D. Gunopulos, "How the live web feels about events," in *CIKM*, Q. He, A. Iyengar, W. Nejdl, J. Pei, and R. Rastogi, Eds. ACM, 2013, pp. 639–648.

[3] J. Parker, Y. Wei, A. Yates, O. Frieder, and N. Goharian, "A framework for detecting public health trends with twitter," in *ASONAM*, 2013, pp. 556–563.

[4] X. Qi, W. Tang, Y. Wu, G. Guo, E. Fuller, and C.-Q. Zhang, "Optimal local community detection in social networks based on density drop of subgraphs," *Pattern Recognition Letters*, vol. 36, pp. 46–53, 2014.

[5] A. Gupta, K. P. Sycara, G. J. Gordon, and A. Hefny, "Exploring friend's influence in cultures in twitter," in *ASONAM*, 2013, pp. 584–591.

[6] C. Bauckhage, "Insights into internet memes." in *ICWSM*, 2011.

[7] J. Leskovec, L. Backstrom, and J. Kleinberg, "Meme-tracking and the dynamics of the news cycle," in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2009, pp. 497–506.

[8] J. Yang and J. Leskovec, "Patterns of temporal variation in online media," in *Proceedings of the fourth ACM international conference on Web search and data mining*. ACM, 2011, pp. 177–186.

[9] J. Teevan, D. Ramage, and M. R. Morris, "# twittersearch: a comparison of microblog search and web search," in *Proceedings of the fourth ACM international conference on Web search and data mining*. ACM, 2011, pp. 35–44.

[10] S. Sen, S. K. Lam, A. M. Rashid, D. Cosley, D. Frankowski, J. Osterhouse, F. M. Harper, and J. Riedl, "Tagging, communities, vocabulary, evolution," in *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*. ACM, 2006, pp. 181–190.

[11] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, "The weka data mining software: an update," *SIGKDD Explor. Newsl.*, vol. 11, no. 1, pp. 10–18, Nov. 2009.

[12] D. W. Aha, D. F. Kibler, and M. K. Albert, "Instance-based learning algorithms," *Machine Learning*, vol. 6, pp. 37–66, 1991.

[13] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.

[14] D. Meyer, F. Leisch, and K. Hornik, "The support vector machine under test," *Neurocomputing*, vol. 55, no. 1a2, pp. 169 – 186, 2003, ¡Support Vector Machines.