Tracing Shifts in Emotions in Streaming Social Network Data

Troels Andreasen[1], Henning Christiansen(✉)[1], and Christian Theil Have[2]

[1] Programming, Logic and Intelligent Systems, Roskilde University, Denmark
[2] Department of Metabolic Genetics, University of Copenhagen, Denmark
Email: {troels,henning}@ruc.dk, christiantheilhave@gmail.com

Abstract. Shifts in emotions towards given topics on social media are often related to momentous real world events, and for the researcher or journalist, such changes may be the first observable sign that something interesting is going on. Further research on why a topic t suddenly has become, say, more or less popular, may involve searching for topics t' whose co-occurrence with t have increased significantly together with the change in emotion. We hypothesize that t' and its increasing relationship to t may relate to a contributing cause why the attitude towards t is changing. A method and tool is presented that monitors a stream of messages, reporting topics with changing emotions and indicating explanations by means of related topics whose increasing occurrence are taken as possible clues of why the change did happen.

1 Introduction

Microblogging through Social media has become a very popular means of communication. Information exchanged is of very diverse nature as is the purposes of posting messages. Messages may report on individuals common day-to-day activities, or be interchanged for chatting. Companies post messages for commercial purposes, while organizations aim to attract attention to their concerns and news media post key headlines to broadcast daily news. It is, however, not always possible to make clear distinctions between types of messages and in many cases there is no need to do so. Patterns of messages and information exchange are, regardless of message types, influenced by events and issues that attract public interest. Often such patterns, as well as changes in these, are good indicators on what's going on and what's considered important issues. Thus analysis of such patterns may in turn lead to expressions of general opinions and trends and may even reveal causes to spotted changes.

In this paper we we describe an approach to querying and monitoring trends, events and opinions in streaming messages. In continuation of the querying approach described in [2] we consider sequences of events and search for sequential patterns with emphasis on changes in attitude. We are aiming at identifying significant changes, that relate to shifts in emotions, sentiments and co-ocurrence patterns. Shifts in emotions towards given topics on social media are often related to momentous real world events. For the researcher or data journalist, such changes may be important clues in their research. Changes can be the first observable sign that something interesting is going on, and a possibility to trace whether this is the case or not can be crucial during research. In our approach a pattern describing, for instance, a shift in emotion can be used in search for interesting cases, while possible causes may be studied among topics with simultaneous shifts in co-occurrence. Further research on why a topic t suddenly has become, say, more or less popular, may involve searching for topics t' whose co-occurrence with t have increased significantly together with the change in emotion. We hypothesize that t' and its increasing relationship to t may indicate a cause why the attitude towards t is changing. One example would be a politician P who looses his credibility from one day the other, while the co-occurrence of P and *corruption* raises from none to very high. Here *corruption* is indicated as a possible reason for P becoming unpopular.

In this study we use twitter data for experiments, where topics on messages are identified with hashtags included in these. Thus the features we take into account for individual messages are, apart from the provided hashtags, timestamps as well as sentiments and emotions derived from the message text by sentiment analysis. The sentiment analysis applied in experiments is provided by a tool described in [7].

The present paper is structured as follows. In section 2 we introduce the language EmoEpisodes and in section 3 we describe the special usage of this language covered here: reporting interesting relationships across topics as indication of possible courses to trends discovered. In section 4 the semantics is described and in section 5 we describe experiments and evaluation based on preliminary implementation of the language. In section 6 we discuss related work and finally in section 7 we conclude.

2 An Emotional Episode Language, EmoEpisodes

EmoEpisodes is a general language for formulating hypotheses about streams of time-stamped messages, concerning properties such as emotions. The language was first introduced in [2] and is in this paper further developed for the purpose of indication of cause. The semantics involves measurements of how well such hypotheses are indicative of or apparent for given time intervals. It can be used as a query language concerning the past and for realtime monitoring with known hypothesis or for discovering new interesting hypotheses.

We assume a continuous time and use $\mathcal{TI}$ to refer the set of all possible time intervals, which may be open or closed. The unit of time is unspecified, but we allow standard units such as days, hours, etc., to characterize intervals. A set of *topics* $\mathcal{T}$ is assumed, e.g., $\mathcal{T} = \{\mathsf{Xmas}, \mathsf{beer}, \mathsf{asteroids}, \ldots\}$, and set of *emotions*, e.g, $\mathcal{E} = \{\mathsf{fear}, \mathsf{happiness}, \mathsf{anger}, \mathsf{sadness}, \ldots\}$. We include sentiment as emotions, $\mathsf{positive}$, $\mathsf{negative}$ and $\mathsf{neutral}$. The *level* of an emotion is characterized by a finite set of symbols $\mathcal{L}$ ordered by magnitude; in this paper we assume the scale $\mathsf{high} > \mathsf{medium} > \mathsf{low}$, but more fine-grained scales may also be used. For the present applications, $\mathcal{T}$ corresponds to twitter hash tags, but it may also be meaningfully extended with other important words extracted from the twitter text.

An atomic emotion *statement* about a topic t associates an emotion E and a level ℓ to that topic, written $t{:}E(\ell)$. For instance "asteroids:fear(high)" measures the degree to which high fear characterizes messages about asteroids. A *data semantics* is given by a *satisfaction degree* function $SD\colon \mathcal{AS} \times \mathcal{TI} \rightarrow [0;1]$ where $\mathcal{AS}$ is the set of atomic statements. It measures how well a given statement characterizes a given time interval. For the present application, $SD(x{:}E(\ell), d)$ reflects the proportion among all messages arriving during time interval d tagged by x that are marked by emotion E; the details are shown in section 4 below.

It is possible for both sentiments positive and negative to be high for a topic t at the same time, if, e.g., 50% of all messages about t are positive and 50% negative. This is quite different from all being neutral. The same goes for intuitively opposite emotions, e.g., love and hate.

An additional atomic statement measures co-occurrences of topics. For topics t and t', the following statement measures the amount of messages the *conditional occurrence* of t' relative to the set of all messages containing t.

- $t'|t{:}(level)$ characterizes the proportion of messages containing t' among those containing t.

As an example: russia|asteroid: (high) measures the degree to which russia characterizes messages about asteroids. Its semantics needs to be defined in a slightly different way than for emotions; details are given in section 4.

EmoEpisodes includes also compound statements that represent an aggregation of atomic statements meant to go for the same time interval; we refer to [2] for details as such statements are not used here.

A *scene* is a statement with an associated *time constraint*; Examples:

asteroids:fear(high)[> 5 days]
asteroids:fear(high)[2013-02-15-15:52:07; 2013-02-16-19:00:00[

The first ones may be applied for different intervals along the time axis, that are longer that 5 days. The detailed language for time constraints is not specified further, and in this paper we need only constraints that fix specific time intervals as shown in the last sample above. we need only constraints that fix specific time intervals as shown in the second sample above.

An *episode* is a sequence of consecutive scenes, separated by semicolon:

asteroids:fear(medium)[2013-02-15;2013-02-17[; doomsday:fear(high)[> 2 days]

The semantics of the full EmoEpisodes language involves first solving the time constraints in order to find a consistent time assignment and then aggregating the *SD* values for each scene. For the present application, we can do with the following extension of *SD* for episodes of two scenes with unique and consecutive time intervals d_1 and d_2; here d denotes the concatenation of the two.

$$SD((s_1[d_1]\,;s_2[d_2]), d) = \min_{i=1,2} SD(s_i[d_i], d_i)$$

Different aggregation operators may be relevant for other applications, as explained in [2], but here, the minimum operator is sufficient.

In fact, EmoEpisodes is a general query language in which variables may stand for unknown constituents, e.g., topics, emotions or degrees, and the query evaluation mechanism may return instantiations of those variables that maximizes the satisfaction degree. An episode S of EmoEpisodes can also be used as a watchdog that signals whenever there is an instance of S and a time assignment ending at the current time, with a satisfaction degree that exceeds a given threshold.

3 Using EmoEpisodes to report interesting, current relationships

Briefly explained, our mechanism uses a watchdog looking for topics t, for which some emotion is changing in a significant way, and then we search for topics t' whose co-occurrences with t are increasing at the same time. Our rationale is an expectation that t' represents an important aspect related to the change in emotion for t. In case the emotions towards such a t' are unchanged during this period, we hypothesize that t' may reflect a cause of the changed attitude towards t; if emotions towards t' changes in a way similar to that of t, we expect the changes for t and t' to have a common cause, and that t and t' may be strongly connected in a symmetric way, perhaps as synonyms for the same thing.

In the following, the constant *now* refers to a present moment during the monitoring of a data stream. To observe a change in an emotion for some topic x, we need to refer to two time consecutive time intervals ending *now* in which the emotion has different values. The constant k refers to the shortest time period that a journalist (or other user) expects an emotion to hold a fairly stable value in order for a following change in that emotion to be significant. An eager sensation oriented journalists will likely prefer a very small k_e, perhaps an hour or 15 minutes (if the complete stream of all tweets from the entire world is available realtime, changes within 15 will be significant). A background journalist may prefer k_e to be 7 days or more. Before the interval $[now - k_e; now]$, a period with a different value for the emotion in question must be observed in order to talk about a change of value. We use a constant k for the length of this pre-period, i.e., $[now - k_e - k; now - k_e[$. The magnitude of k should depend on the sort of phenomena of interest (are they normally stable or fluctuating) combined with considerations about the overall arrival frequency of messages.

When a change in emotional value has been observed, we look for topics t' whose co-occurrences with t are increasing. As above, we define a similar constant k_o and reuse k for two similar, successive intervals. The rationale for having two different constants k_e and k_o is an expectation that when a (perhaps drastic) real world event changes the general attitude towards a topic t, it may take some time before more interesting circumstances or speculations about the event becomes known and discussed; this can motivate $k_o < k_e$. Different values of k may be used for measurements of t and t', but we see no good reason for that. We expect identical values for k_e and k_o to be acceptable, although more empirical testing may be needed to find the best choice.

Our system is based on the following abstract algorithm; T_0 is a subset of all topics declared by the user to be of interest. A threshold $\theta \in [0;1]$ for significant degree of satisfaction is assumed.

1. Identify the set of all topics $t \in X \subseteq T_0$ for which there exist emotion E and levels $\ell_0 \neq \ell_1$ such that the satisfaction degree of the following query is $\geq \theta$,

$$t{:}E(\ell_0)[now - k_e - k; now - k_e[\ \mathbf{;}\ t{:}E(\ell_1)[now - k_e; now].$$

2. For each $t \in X$, identify the set of $t' \in Y_t$ for which there exist levels $\ell_2 < \ell_3$ such that the satisfaction degree of the following query is $\geq \theta$,

$$t'|t{:}(\ell_2)[now - k_o - k; now - k_o[\ \mathbf{;}\ t'|t{:}(\ell_3)[now - k_o; now].$$

Notice that the sets Y_t need not be subsets of T_0.

The algorithm may run continuously over time, as *now* inevitably moves forward, although an implementation may need to use a more or less fine-grained discretized time in order to reduce the computational overhead. The results from this algorithm is monitored, showing for each t in the current X set,

- the list of emotions E that may trigger step 1 above,
- the list Y_t with, for each $t' \in Y_t$, whether t' is a *possible cause* of the emotion change (i.e., emotions towards t' stable), or a *related concept* (i.e., emotions change similarly to t); measured analogously to step 1 above.

A user interface to be used in, say, an editorial office may display this rudimentary information as an effective way for the journalists quickly to recognize potentially new hot topics; an proposal for this is shown in section 5 below. Additional information may be called up by click buttons, e.g., about actual satisfaction degrees and detailed information about the t' concepts.

4 A Data Semantics for Mining Trends on Twitter

Different choices for data semantics and aggregation operations are introduced in [2] for a variety of applications and views of the data. Here we give the data semantics used for the present application of measuring changes in streaming twitter messages (which coincides with the so-called elitist semantics of [2].

As indicated above, we include the so-called sentiment as a special sort of emotion. Let us make this precise,

$$\begin{aligned}
\mathcal{E} &= \mathcal{E}_1 \cup \mathcal{E}_2, \quad \text{where}\\
\mathcal{E}_1 &= \{\mathsf{anger}, \mathsf{disgust}, \mathsf{fear}, \mathsf{joy}, \mathsf{sadness}, \mathsf{surprise}\}\\
\mathcal{E}_2 &= \{\mathsf{negative}, \mathsf{neutral}, \mathsf{positive}\}
\end{aligned}$$

When a twitter message arrives, it is classified by a filter with an emotion $\in \mathcal{E}_1$ and a sentiment $\in \mathcal{E}_2$; however, the filter may fail in identifying emotion or sentiment. The same message cannot be classified with two different proper emotions or two different sentiments. We introduce some notation; we assume topic $t \in \mathcal{T}$, time interval $d \in \mathcal{TI}$, $E \in \mathcal{E}$.

$\delta_\sharp(t,d)$: The set of messages tagged by topic t during d.

$\delta_i(t{:}E,d)$: The set of messages tagged by topic t during d and classified by the filter by $E \in \mathcal{E}_i$.

$\delta_i(t,d)$: The set of messages tagged by topic t during d and classified by the filter by some $E' \in \mathcal{E}_i$.

The relative frequency of emotion or sentiment $E \in \mathcal{E}_i$ for a topic t during time interval d is defined as

$$R_i(t{:}E,d) = \frac{|\delta_i(t{:}E,d)|}{|\delta_i(t,d)|}.$$

For topics t, t', we define the relative frequency of t' given t during d as

$$R_\sharp(t'|t{:},d) = \frac{|\delta_\sharp(t',d) \cap \delta_\sharp(t,d)|}{|\delta_\sharp(t,d)|}.$$

The different levels $\mathsf{high} > \mathsf{medium} > \mathsf{low}$ are treated as simple fuzzy linguistic terms in the definition of satisfaction degrees for atomic statements, using membership functions μ_i^ℓ for $i \in \{1,2,\sharp\}$ and level $\ell \in \{\mathsf{high},\mathsf{medium},\mathsf{low}\}$. The satisfaction degrees defining the data semantic are now given as

$$SD(\phi_i(\ell),d) = \mu_i^\ell(R_i(\phi_i,d))$$

where ϕ_i is one of $\phi_1 = (t{:}E_1)$, $\phi_2 = (t{:}E_2)$, $\phi_\sharp = (t'|t{:})$. Figure 1(a) shows definitions of relative satisfaction level terms high, medium, low for classifications $\mathcal{E}_1$ and figure 1(b) for classifier $\mathcal{E}_2$ (notice that $|\mathcal{E}_1| = 6$ and $|\mathcal{E}_2| = 3$). For occurrences of t' given t, we choose (arbitrarily) the relative frequency $f_{t'}$ of t' among all tweets, measured for a sufficiently large period D in the past, as the midpoint for medium; the remaining membership functions are defined accordingly.

$$f_{t'} = \frac{|\delta_\sharp(t,D)|}{|\delta_\sharp(D)|} \quad \text{where } \delta_\sharp(D) \text{ is the set of tweets arriving during } D.$$

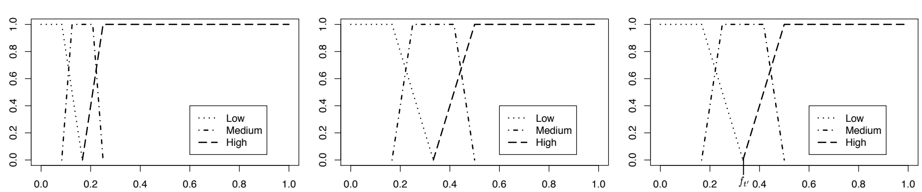

Fig. 1. Membership functions μ_ℓ for fuzzy linguistic terms over relative satisfaction with level $\ell \in \{\mathsf{high},\mathsf{medium},\mathsf{low}\}$. Shown in (a) for classifier $\mathcal{E}_1$, in (b) for classifier $\mathcal{E}_2$ (6 and 3 classes respectively), and in (c) for $t'|t$: that depends on t' but not on t.

5 Experiments

Our system is not fully developed with a finished user interface for reporting the results. To illustrate the utility of the mechanism described in the paper we show a manually created report for results generated from a sample run of the system on a large corpus of tweets extracted from the Twitter firehose during the period from December 23 2012 to February 7 2013. The *firehose* is an API which gives access to (a random fraction of) tweets as they occur in real-time. The corpus thus represents a random sample of tweets which can be seen as representative of all tweets in the period.

We use this corpus as if we had continuously been collecting tweets and the current date – *now* – happens to be one of the days up to February 7 2013. We define time constraints in a granularity of days; We use a setting of $k = k_e = k_o = 5\ days$.

After filtering uninteresting nuisance tags such as *teamfollow, followback,* ... which are used in pyramid-schemes to gain followers and retweets, the best scoring result involves the tag *iphone* which co-occurs with the tag *gameinsight*.

In the result shown in the report in figure 2 we see that the system has detected an increase in negative sentiment for the topic *iphone* and that it is correlated with an increased co-occurrence of the tag *gameinsight*. This is a tag which is mostly used by certain games when they post automatic updates to Twitter and promote a viral effect with tweets such as,

– "just reached level 19 on Rock the Vegas on my iPhone http://t.co/Z4nLdor1 ♯iphone iphonegames ♯gameinsight"
– "I have completed the quest 'Order 3 Long-Term D...' in the ♯iPhone game The Tribez. http://t.co/nn6tb60a ♯iphonegames, ♯gameinsight"

Of the above tweets, which are announcements of game progress, the second one is classified as having negative sentiment mainly because of the word *quest* which occur in the negative sentiment classifier dictionary. When the word occurs, it often results in negative sentiment classification of the tweet in which it occurs. While in the context of this tweet the classification seems dubious, a sudden surge in popularity in *The Tribez* game, which publishes a lot of tweets about *quests*, gives rise to an increased negative sentiment about *iphone* as can be observed figure 2.

6 Related work

As mentioned above the approach described in this paper is a continuation the querying approach described in [2]. While [2] took a more general approach towards a language for specifying patterns the present paper has a primary focus on identification of cooccurring topics that may candidate as possible causes.

Our terminology is inspired by the seminal work of [8] who suggested a way to define episodes in sequences of discrete events (from a finite alphabet of such) and gave algorithms to search for a sort of association rules among such episodes.

Hot topic: iphone **Changed emotions:** negative-sentiment↑

Increased co-occurring topics

Possible causes: gameinsight

Related: app,iphonegames, Apple, iPad, ringtones

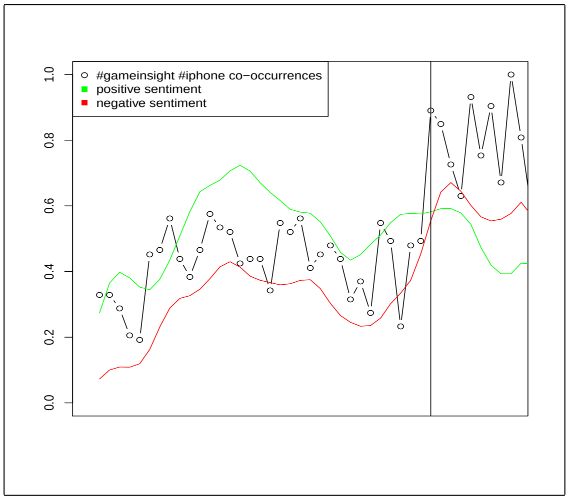

Fig. 2. The figure shows a sample report for the *"hot topic"* result *iphone*. The associated emotion changes are indicated and the most promising candidate is the co-occurring tags, (here only *gameinsight*), are listed. Similarly, *related* highly co-occuring topics which do not correlate sufficiently with the emotion changes are listed. The report also visually plots the change in both emotion (negative sentiment) measured as number of emotion classifications for the topic (*iphone*) scaled to a unit interval and the number co-occurrences also scaled to a unit interval for comparability. For reference, we also provide a similar plot for positive sentiment change (which did not result in an alert).

Before that, [1] described algorithms for mining frequent, sequential patterns in a transaction database. See a recent survey [10] on later work inspired by [1, 8]. Our work differs in that we present a logical language for specifying scene and episode sequences that refer to measurements over large sets of timestamped tweets, rather than over a finite alphabet. The language has a well-defined, graduated truth semantics, that by parameterization allows different interpretations of the data.

Analysing trends in blogging corpora based on sentiment is attracting increasing interest by researchers and the micro blogging platform Twitter have shown to be a valuable ressource for this purpose in part due to users widespread use of hashtags that to some extent can be interpreted as topics. In [12] the authors analyse trending topics with emphasis on duration in an attempt to distinguish major public events. In this approach so-called provenance of topics plays a key role. The goal is the same a ours – to search for important topics and trends related to these. However, while we analyse patterns of attention and opinion during consecutive time intervals, the approach in [12] is to apply and additional resource as context, Wikipedia, in an attempt to distinguish important topics and significants trends, such as those related to major public events. Another approach to identification of trends with similarities to ours is described in [9]. They consider real time trend detection over the Twitter stream and based on this provide a monitor system. While we attempt to identify trends that range over changes of attitudes, [9] identifies trends with emerging topics on Twitter, and aims to synthesize accurate descriptions of topics.

An approach that combine sentiment analysis and volume-based measures like ours is described in [4]. They investigate the potential to model political sentiment through mining of social media, and they indicate with their results that their combined approach to analysis may provide prediction for their case, an election campaign, when including examination of sample sizes, time periods as well as methods for qualitatively exploring the underlying content. Further in [3] the authors argues that analysis on the short document length in microblogs provide compact and explicit sentiments. They argue that it is easier to classify the sentiment in short form documents than in longer documents.

Sentiment analysis of Twitter messages over time has previously been demonstrated to correlate with public opinion measured by Gallup polls [11]. However, in this study only positive/negative sentiment is measured and they do not provide a method to search for specific patterns. More detailed emotions are used in a study which correlates Twitter sentiments to socio-economic phenomena [6]. This group have also studied correlation of sentiment of Twitter messages to stock prices [5] and also find sentiment to be highly correlated with stock prices, but their approach does not specifically consider surprising events.

7 Conclusion and future work

In this paper we have presented a language for querying development of attitudes towards given topics over time and we have illustrated its use with Twitter data. Expressions in the language may include specification of topics, and emotions and sentiments related to these may be queried within specified time intervals. The general form of a query specifies two consecutive varying length time intervals that can be considered moving along the time line in search for significant changes from the one to the next interval. A query can be open wrt topics as well as wrt emotions. Thus we can look for changed attitude towards specific topics or in general search for topics characterised by such changes. We can also search for specific attitude changes, such as increased surprise, or in general search for significant changes along any emotion or sentiment in any direction. We have put a special emphasis on mining causes. Thus we are not only concerned with significant changes, but also, as part of the framework, aimed at deriving indications of causes by investigating topic co-occurrences.

Being able to answer queries, spot changes and indicate causes, as supported by our language, can be very useful for journalists and researchers looking for interesting new trends and it may have important implications for, e.g., social, socio-economical, political science and for market analytics. Analysis of Social media data and mining for trends is not new, but to our knowledge, a general approach, introducing a query language, and a support for mining, not only for changes, but also for indications of possible causes, has not been seen before.

Our prototype implementation is preliminary and can be improved in a number of ways. Firstly the language is not fully implemented and currently some of the data analytics has to be initiated manually. Secondly the mentioned "watchdog" use of the approach – supporting in principle any number of query expressions introduced as search agents for automated notifications – has not been implemented yet and decisions has to be made concerning its specific functionality.

References

1. R. Agrawal and R. Srikant. Mining sequential patterns. In P. S. Yu and A. L. P. Chen, editors, *ICDE*, pages 3–14. IEEE Computer Society, 1995.
2. T. Andreasen, H. Christiansen, and C. T. Have. Querying sentiment development over time. In H. L. Larsen, M. J. Martín-Bautista, M. A. Vila, T. Andreasen, and H. Christiansen, editors, *Flexible Query Answering Systems - 10th International Conference, FQAS 2013, Granada, Spain. Proceedings*, volume 8132 of *Lecture Notes in Computer Science*, pages 613–624. Springer, 2013.
3. A. Bermingham and A. F. Smeaton. Classifying sentiment in microblogs: Is brevity an advantage? In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*, CIKM '10, pages 1833–1836, New York, NY, USA, 2010. ACM.
4. A. Bermingham and A. F. Smeaton. On using twitter to monitor political sentiment and predict election results. In *Proceedings of the IJCNLP conference*, Chiang Mai, Thailand, 2011.
5. J. Bollen, H. Mao, and X. Zeng. Twitter mood predicts the stock market. *Journal of Computational Science*, 2(1):1–8, 2011.
6. J. Bollen, A. Pepe, and H. Mao. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*, pages 450–453, 2011.
7. T. P. Jurka. sentiment: Tools for Sentiment Analysis. Version 0.2. http://github.com/timjurka/sentiment.
8. H. Mannila, H. Toivonen, and A. I. Verkamo. Discovery of frequent episodes in event sequences. *Data Min. Knowl. Discov.*, 1(3):259–289, 1997.
9. M. Mathioudakis and N. Koudas. Twittermonitor: Trend detection over the twitter stream. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*, SIGMOD '10, pages 1155–1158, New York, NY, USA, 2010. ACM.
10. C. H. Mooney and J. F. Roddick. Sequential pattern mining – approaches and algorithms. *ACM Comput. Surv.*, 45(2):19:1–19:39, Mar. 2013.
11. B. OConnor, R. Balasubramanyan, B. R. Routledge, and N. A. Smith. From tweets to polls: Linking text sentiment to public opinion time series. In *Proceedings of the International AAAI Conference on Weblogs and Social Media*, pages 122–129, 2010.
12. T. Tran, M. Georgescu, X. Zhu, and N. Kanhabua. Analysing the duration of trending topics in twitter using wikipedia. In F. Menczer, J. Hendler, W. H. Dutton, M. Strohmaier, C. Cattuto, and E. T. Meyer, editors, *ACM Web Science Conference, WebSci '14, Bloomington, IN, USA*, pages 251–252. ACM, 2014.