

Dynamics of News Events and Social Media Reaction

Mikalai Tsytsarau
University of Trento
tsytsarau@disi.unitn.eu

Themis Palpanas
Paris Descartes University
themis@mi.parisdescartes.fr

Malu Castellanos
Hewlett Packard
malu.castellanos@hp.com

ABSTRACT

The analysis of social sentiment expressed on the Web is becoming increasingly relevant to a variety of applications, and it is important to understand the underlying mechanisms which drive the evolution of sentiments in one way or another, in order to be able to predict these changes in the future. In this paper, we study the dynamics of news events and their relation to changes of sentiment expressed on relevant topics. We propose a novel framework, which models the behavior of news and social media in response to events as a convolution between event’s importance and media response function, specific to media and event type. This framework is suitable for detecting time and duration of events, as well as their impact and dynamics, from time series of publication volume. These data can greatly enhance events analysis; for instance, they can help distinguish important events from unimportant, or predict sentiment and stock market shifts. As an example of such application, we extracted news events for a variety of topics and then correlated this data with the corresponding sentiment time series, revealing the connection between sentiment shifts and event dynamics.

1. INTRODUCTION

The problem of monitoring the evolution of sentiment has been studied in the context of different research areas, from social studies to reputation management [13]. However, there is still a lack of understanding of what causes the crowd’s sentiment to change. Some people change their attitude for personal reasons, some do so influenced by their connections, but most likely people change their opinion when a new evidence comes into their consideration.

By aggregating sentiments, expressed in multiple texts, and assessing the result with statistical measurements, we can capture certain changes, or shifts, in global sentiment, which cannot be attributed to random variation [12]. Understanding the reasons for these sentiment shifts provides insight about products, services and policies to companies and institutions. Thus, this problem has gained a lot of attention in academia and industry [16, 10, 4, 7].

In this paper, we propose studying the correlation of important shifts in sentiment to news events reported by different news sources. The objective of this paper is to exploit the current work on senti-

ment analysis and contradiction detection, and extend the state of the art by understanding news events dynamics and modeling relationships *between sentiment changes and news events*.

We observe that most news events are announced as atomic pieces of information and their impact is not readily intelligible from the text alone. To determine the importance and impact of news to people, it is crucial to consider the publication dynamics of the whole crowd, rather than only from news agencies or news media [10]. Analyzing the aggregated publication volume or user interest on a specific topic over time can help in understanding the event’s importance and dynamics. Earlier detection of such events can be useful to publishers, content providers and search engines [9]. However, in certain cases social media can contribute to this volume all by themselves (without any external stimuli, such as a news event coming from a news agency) and also maintain a trending volume growth over long time periods. These effects distract the observed events dynamics and may even make them undetectable.

Our method addresses these problems by representing publication dynamics as the result of the interplay between the original news’ importance and social response. More specifically, our modeling is based on the idea that news media (or social media) can be described by a special “response” function, which determines the resulting dynamics of news publication (or user interest) for an event. This opens up the possibility of recovering the original event sequence and its varying importance over the time dimension.

The response function can be seen as a likelihood of a “delayed” publication for an event: similar to a phone conversation, where a

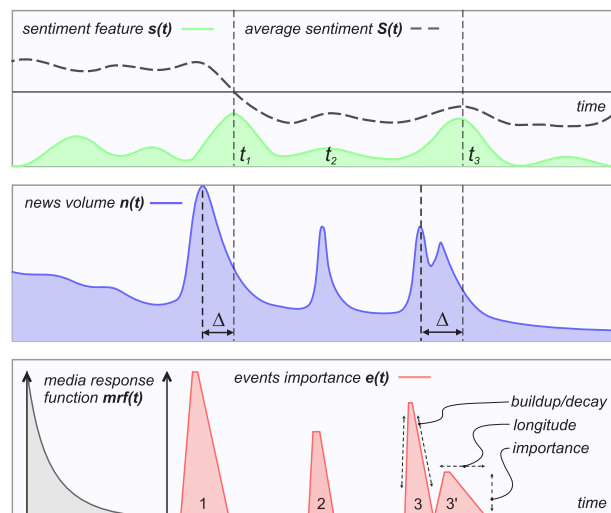


Figure 1: An example time series, with sentiment shifts (top) correlated to news volume (mid) and events importance (bottom).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
KDD'14, August 24–27, 2014, New York, NY, USA.
Copyright 2014 ACM 978-1-4503-2956-9/14/08 ...\$15.00.
<http://dx.doi.org/10.1145/2623330.2623670>.




Figure 2: Some of the events identified by our system for “iPad”.

delay in circuits creates the “echo” effect, news and social media tend to re-publish, cite, and discuss previous articles, extending this activity over time. In this case, the peak intensity of publications does not always coincide with the beginning of the event, or its peak importance. To tackle these problems and recover the original event importance, we employ deconvolution, which is a popular technique for improving audio or image quality [3].

Based on this method, we are able to accurately measure events importance and response dynamics, opening up the possibility to evaluate their impact on sentiment evolution. The results obtained by our study indicate the existence of different types of response dynamics in Twitter, and reveal their impact on sentiment changes.

1.1 Motivating Scenarios and Examples

We would like to make the reader more familiar with our goals by introducing the following use cases.

Example 1: Consider a search interest [9] time series extracted from Google for the topic “iPad”, shown in Figure 2, blue. It features a growing number of search queries overlaid with a series of overlapping bursts of user interest, making it very hard to detect news events. The output time series (Figure 2, red), processed using our method, demonstrates a more vivid event separation and dynamics, making them easily detectable. Moreover, the output time series appears without the global trend, which is processed alongside the deconvolution. We note, that achieving the same effect is not possible by subtracting a moving average or by using linear regression, since the presence of peaks in time series biases the trend estimation.

Example 2: It is known, that response dynamics is often determined by events, and can be used to classify events by their impact and expectedness [2, 5]. For example, in the time series shown in Figure 3, we see that the death of Michael Jackson was so impacting to media, that their response lasted for over a year. On the contrary, the interest to Harry Potter movies has usually a two-month period of decay. Moreover, we observe that these events have different decay patterns (we further discuss this point later in the paper). Our method has been designed to detect and analyze different response dynamics, making suitable for cases like the above.

Example 3: Since its early announcement, Samsung’s “Galaxy Tab” was regarded in social media as a very fine competitor to Apple’s “iPad”, receiving mostly positive sentiments. However, the attitude of people dramatically changed to negative at the moment when Samsung published the price for the device. By observing the media reaction to these and other impacting news, it becomes possible to predict changes in their opinion more rapidly, i.e., as soon as we are able to recognize the establishing trend dynamics.

$S(t)$	Raw sentiment time series	$\xi(t)$	Random variable
$s(t)$	Sentiment feature time series	$\rho(\cdot)$	Correlation function
$e(t)$	Event importance time series	$\delta(t)$	Dirac (pulse) function
$n(t)$	News volume time series	$h(t)$	Heaviside (step) function
$x(t)$	Aggregated (total) news volume	Δ, δ	Lags between time series
$rf(t)$	Response function of a user	$\gamma(t)$	Importance decay function
$mrf(t)$	Response function of entire media	τ_0, α	Response decay parameters

Table 1: Notations used in this paper.




Figure 3: Dynamics of Google Search volume for different topics. Extracted event importance is shown as triangles (normalized).

1.2 Contributions

To the best of our knowledge, this is the first work that proposes a principled modeling of news dynamics in various media and news interaction with sentiments. Our main contributions are summarized as follows.

- We analyze the existing models of publication dynamics, and discuss their foundational principles.
- We model the dynamics of news volume as a convolution between events importance and media response function.
- We develop a method for news events extraction, based on deconvolution with automated parameter optimization.
- We assess several sentiment features for their correlation to news events, and the possibility of predicting their changes.

Since our framework relies on deconvolution, it can accommodate various response functions, suitable for different cases. We note that our method does not require describing the news publication dynamics using a differential equation. Moreover, it is possible to automatically learn media response functions from the data.

2. FORMALISM

We summarize the most important notations used in our paper in Table 1 and discuss them in Section 2.1. Following that, we introduce our problems in Section 2.2.

2.1 Definitions

We are given a time series of numeric values, $s(t)$, which is derived from raw sentiments $S(t)$ for a particular topic T and represents some sentiment feature or interestingness measure, e.g., contradiction level [15]. These time series are shown in Figure 1 (top), where $s(t)$ represents the contradiction level of raw sentiments $S(t)$. Along with the sentiment time series, we are given the news frequency time series represented by $n(t)$ (Figure 1, middle), and the corresponding correlation function $\rho(s, n)$, which takes both time series as input parameters and computes a real-valued correlation coefficient. We note that $\rho(s, n)$ can be a special function adapted for particular measures used to compute $s(t)$ and $n(t)$. For instance, $n(t)$ can be a rather bursty time series, where the notions of mean value and deviations from the mean make no apparent physical sense. Thus, it can be more effective in this case to consider a correlation measure based on derivatives, or on bursts alignment. Finally, we use the event importance time series $e(t)$ to model event dynamics, and response functions $rf(t)$ and $mrf(t)$ (represented in Figure 1, bottom) to model social and news media behavior. We describe these functions in more detail in Section 3.

2.2 Problems

We are given a collection of sentiment feature time series $s(t)$ and news volume time series $n(t)$ for various topics. For a particular topic T , we want to detect the most important sentiment shifts according to $s(t)$, and identify the news events, which have caused them, so that it will be possible to predict future sentiment changes on that topic upon observing and analyzing relevant news events. This general problem can be decomposed into a set of sub-problems with the scopes of news-sentiment interaction and news-events interaction respectively.

News and sentiments interaction can be coarsely modeled by correlation. However, the problem of finding dependencies between sentiments and news is inherently complex, since sentiment shifts may be correlated not only with the news on the same topic, but also with several other (related) topics.

PROBLEM 1. Given $s(t)$ for a topic T and a correlation measure $\rho(s, n)$, determine a list of topics T_i , whose news time series $n_i(t)$ correlate with $s(t)$. For each correlation pair, determine a time lag between the two series.

After determining a substantial amount of news events, which caused sentiment changes, it may be possible to predict shifts in sentiment for related topics by extracting features of the recognized event and using them as input to a classifier model. Related topics can be determined by referring to the list of news time series, which are correlated with the given sentiment time series. The classifier model can be trained on a dataset of confirmed causality cases between news events and sentiment shifts. Alternatively, it can use automatically extracted correlation pairs. We note that this problem requires a deeper study with complex causality modeling, and we defer it to a future study. In this paper, our goal is to explore the event dynamics in different media, and establish the connection between events and sentiment shifts.

3. BACKGROUND AND RELATED WORK

There exist approaches that tackle specific aspects of our problem in various domains, yet, they cannot be combined to solve the problem that our approach solves. In this section, we provide a detailed evaluation of their properties and design principles, and use it to establish a theoretical and empirical background for our method.

To the best of our knowledge, there is currently no other system, or method, capable of automatically detecting impacting events and identifying properties of their dynamics from news publication volume. There exist only few news tracking methods which are able to recognize different event types, but neither of them provides numeric characteristics useful for event modeling and sentiment prediction purposes, such as a correct time framing and importance level of an event. Still, they are not capable of distinguishing adjacent events when their publications are highly overlapping. Finally, the existing dynamics models target only individual news propagation media (e.g., blogosphere, microblogging, youtube), and have not been generalized (or tested) on multiple domains.

3.1 Event Models

Lehmann et al. [5] study collective attention in Twitter and its propagation through user networks. They measure the aggregated volume before, during and after the event's peak by subtracting the baseline level of attention (computed using a sliding window). Based on the relationship between these three values, they define four classes of news events according to their expectedness and impact. Accordingly, events can be: a) *expected impacting*, where there is a growth of volume before an event (anticipation) and a

decay afterwards (response); b) *expected non-impacting*, where the event's outcome is of a lesser concern than the event itself; c) *unexpected impacting*, featuring an instant appearance and a lengthy response; d) *unexpected non-impacting* or *transient*, where neither the event nor its outcome are important to the media;

Alternatively, Crane and Sornette [2] consider events as having internal (*endogenous*), or external (*exogenous*) origin, and being of either *critical*, or *sub-critical* importance to social media. We observe that *exogenous critical* and *exogenous sub-critical* event types in their classification coincide with *unexpected impacting* and *unexpected non-impacting* event types from [5]. In addition, the authors also introduce a concept of *endogenous* events, which is broadly similar to memes in social media.

While these classes capture events semantics on a very good level, a more detailed analysis of news dynamics requires reliable extraction of peak shapes and proper modeling of social media. The first requirement is important since publication volume often contains noise and (evolving) background level, which masks individual peaks. The second requirement is necessary to distinguish between the volume generated by endogenous factors (imitation) and that generated by exogenous factors (news importance). Therefore, we investigate the most recent models of news dynamics and their ability to correctly describe and predict news volume.

3.2 Meme Model

Memes in social media are the outbursts of publications on some topic, which can be assigned to the "endogenous critical" type of the above classification. Since they have no particular external driving force or impact on sentiment, we want to distinguish them from "expected impacting" events. The topic of news dynamics in social media (and memes in particular) was studied by Leskovec et al. [6], who propose a model of meme and news dynamics based on three assumptions for the interaction of news sources: *imitation*, *recency preference* and *concurrency*. The imitation (endogenous) hypothesis assumes that news sources are more likely to publish on events that have already seen large volume of publications. The recency hypothesis marks the tendency to publish more on recent events. Finally, the concurrency hypothesis states that news sources have a limited capacity and choose only one event at a time to report on. The authors express the imitation and recency preferences through the functions of a combined news volume, $f(x)$, and time passed since the beginning, $\gamma(t)$:

$$n(t) = \frac{dx}{dt} = c \cdot f(x)\gamma(t) \quad (1)$$

In the above equation, dx/dt represents the amount of publications at time t , and x represents the total amount integrated since the beginning. We note that whereas equation (1) is based on the assumptions given in [6], it is different to the equation studied in that paper, where $x + dx/dt$ were mistakenly used as $n(t)$ ¹.

Leskovec et al. demonstrate that in the simulated environment consisting of concurrent news sources, both time and volume components are necessary to generate the oscillating nature of news volume. However, repeating this experiment with the original meme time series revealed a substantial deviation between predicted and observed volumes, when the same assumptions (namely, global time and volume) were used. First, the global volume (accumulated since the beginning of the time series) is not useful to predict any but the very first peak. Second, the global time (either from the beginning of the time series or from the first peak) is both arbitrary (which peak is the first? when a time series begins?) and can not capture subsequent peaks. This said, we assume that the authors

¹J. Leskovec. Personal communication, June 2013.

implied the global nature of their model with respect to parameters, while the time and the volume were local for every seeded event. Nevertheless, we believe that equation (1) intuitively captures some of the properties of news dynamics, and is therefore worth considering.

3.3 Stochastic Multiplicative Model

Asur et al. [1] propose a model for news dynamics, described by stochastic multiplicative process, driven by independent random variables (noise) and a time-decaying variable (recency). Based on this model, they predict a linear initial growth of publications volume, and a log-normal distribution of this volume over different time series. Both hypotheses are supported by an empirical evaluation on Twitter [1] and Digg [17] data. However, this model is not useful for prediction purposes, since it relies on a mixture of random variables (located at subsequent time intervals). Even though it is possible to infer values of these variables for past time intervals, the hypothesis that they are independently distributed forbids estimation of subsequent variables.

A closer look at this model reveals that it is formulated in a recursive manner, where the total volume accumulated by a time tick $t + dt$ is expressed through the volume at time t , i.i.d. random variable $\xi(t)$ and time-decaying component $\gamma(t)$:

$$x + dx = [1 + \gamma(t)\xi(t)]x \quad (2)$$

To analyze the proposed dynamics, the above equation can be transformed into a more convenient form (remember though, that dx/dt , x and t remain discrete):

$$n(t) = \frac{dx}{dt} = \gamma(t)\xi(t)x \quad (3)$$

Now it can be clearly seen that the volume of news published at time t depends on the previously accumulated volume x , discounted by $\gamma(t)$, and on random variable $\xi(t)$. Although the dependency of (3) on random variables forbids its analytical derivation, we observe that it is very similar to (1), and can also be represented as the product of exponent and time factors, especially in the case when $\xi(t)$ are not i.i.d.

Another important observation in favor of the continuity of $\xi(t)$ is that these variables may represent an exogenous factor, the news importance $e(t)$, which pushes volume up and counteracts the decreasing trend of $\gamma(t)$. Comparing (1) and (3) under this perspective, it becomes evident that they are essentially the specific cases of a more general model, multiplying endogenous and exogenous factors. Our model differs from these two by considering a convolution between endogenous and exogenous factors.

3.4 Hawkes Poisson Process Model

A study of social system's dynamics by Crane and Sornette [2] comes the closest to our work with regard to a modeling based on user response. The authors study dynamics of book sales [8] and social content [2], based on a widespread model of hyperbolic (long-memory) user response function $rf(t) \sim 1/t^{1+\theta}$, $0 < \theta < 1$. Taking the ensemble average of a Hawkes Poisson Process driven by this response function and a spontaneous rate $e(t)$, they express $n(t)$ conditional on itself and on an average branching ratio μ :

$$n(t) = e(t) + \mu \int_{-\infty}^t rf(t - \tau)n(\tau)d\tau \quad (4)$$

Correspondingly, equation (4) takes the form of the convolution between event importance $e(t)$ and a media response kernel $mrf(t)$, considered as the output of (4) in the case when $e(t) = \delta(t)$:




Figure 4: Interaction between key components of the system.

$$n(t) = \int_{-\infty}^t mrf(t - \tau)e(\tau)d\tau \quad (5)$$

Summarizing the above studies, we see that the imitation factor does not play as important role in news dynamics as in meme dynamics. Moreover, [1] and [5] observe that propagation of news through user networks is not epidemic, i.e., it rather depends on news importance than on numbers of followers of users who spread the news. Therefore, we consider publication likelihood being dependent more on recent volume than on past volume, deviating from the purely endogenous model. In addition, we consider that news events have continuity and varying importance, which also affect the publishing dynamics (exogenous assumption). These assumptions require a more complex modeling of news dynamics, but result in more accurate models, which we discuss in the following sections.

4. PROPOSED METHOD

Our task requires processing two different kinds of data, i.e., sentiments and news volume, which come from different sources and also at very different rates. In Figure 4, we show the overall architecture of our system by dividing its major components into two main layers: *News Layer* and *Sentiment Layer*.

In the *News* layer, we aggregate the volume of the news for a topic into a time series, which will be further analyzed to detect news events. To detect a news event and extract its features we perform a deconvolution of news volume time series (or its fragment).

The *Sentiments* layer aggregates sentiments over time for a specific topic, and detects interesting changes, which can be contradictions, outbursts of sentiments' volume, or other situations.

4.1 Detecting Impacting Events

As we already noted, not every kind of publications outbursts is caused by external news. For instance, "endogenous critical" and "expected impacting" events may produce a similar response in social media, yet only the latter one is relevant to our study. Moreover, not every kind of news dynamics has an impact on sentiment, so we want to distinguish them at a fine level of detail.

Following these observations, we introduce our model for social media and news dynamics. We start with the description of basic social media responses and our representation of event importance. Then, we introduce a novel method to extract important properties of events, which is based on deconvolution, and propose methods to estimate the parameters for this process.

4.1.1 Modeling News Dynamics

We model the observed news dynamics (frequency of publications) as a response of social media to external stimuli. We represent the output as a convolution of two functions: the news events importance sequence and the media response function:

$$n(t) = \int_{-\infty}^{+\infty} mrf(\tau) \cdot e(t - \tau)d\tau \quad (6)$$

where $mrf(t)$ is the media response function (in general, decaying), and $e(t)$ is the actual event sequence, which is unobserved. We depict these functions in Figure 1, bottom.

However, in order to recover the original event sequence, we need to perform a deconvolution of news frequency time series, for which we should know the exact shape of $mrf(t)$. To model the behavior of media, we propose using a family of normalized decaying functions, which have the aggregated volume equal to 1.0 and which are defined on $t > 0$ using the Heaviside step function $h(t)$. For instance, these functions can be *linear*, *hyperbolic* or *exponential*, as demonstrated in Figure 5 and formalized below.

$$\text{linear} \quad mrf(t) = \left(\frac{2}{\tau_0} - \frac{2t}{\tau_0^2} \right) h(t)h(\tau_0 - t) \quad (7)$$

$$\text{hyperbolic} \quad mrf(t) = h(t) \frac{\alpha-1}{\tau_0} \left(\frac{t+\tau_0}{\tau_0} \right)^{-\alpha} \quad (8)$$

$$\text{exponential} \quad mrf(t) = \frac{1}{\tau_0} e^{-t/\tau_0} h(t) \quad (9)$$

We note, that linear response has the shortest effect and hyperbolic response has the longest effect on time series if we consider comparable values of decay time τ_0 , which takes on a role similar to the half-life of radiation decay.

We consider decaying response functions, because news events become obsolete and cease being published very soon after their appearance. Another possible explanation for this phenomenon is the saturation of content providers: the likelihood (the temporal rate) of news publication in the absence of events activity is inversely dependent on the number of news, which have been previously published. Moreover, the shapes of response functions convey additional information regarding impact and expectability of events, and also reveal properties of the media.

Linear Response (7). In the case of a linear response, the probability of publishing on an event linearly decreases with time, and the media cease publishing after a finite cutoff time τ_0 . Linear response is characterized by a constant rate of content generation, and results in nearly linear dynamics of news volume for spike event shapes. We are interested in evaluating it since this kind of dynamics was observed by [1] for event buildups.

Exponential Response (9). For the exponential response, the rate of decline is proportional to the current probability value at any moment, hence the decay is initially more rapid than the linear, but becomes less pronounced towards the end, with the probability reaching 0 in an infinite time.

Hyperbolic Response (8). In the case of hyperbolic (power law) response, the probability follows a more pronounced decay compared to exponential (the decay rate is usually quadratic to the current value). It decreases very rapidly when the time τ is small, but has a heavy tail afterwards. Here, parameters $\alpha > 1$ and $\tau_0 > 0$ control the sharpness of a response. The hyperbolic response is very interesting since its amplitude can reach infinity in a constant time (unlike the exponential) and in the case of $\alpha \approx 1$ its rate of decline is proportional to the square of its current value.

Symmetric and Asymmetric Responses. Aside from being impacting or non-impacting, events can be anticipated or unexpected. These properties of events manifest themselves through the observed response of the media. For instance, unexpected events will have no buildup of volume preceding the event (consider earthquakes), and expected non-impacting events will have abrupt decay of volume upon their end (consider football finals). Therefore,

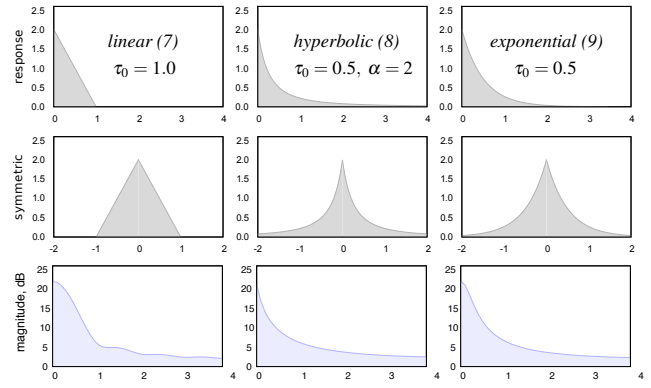


Figure 5: Media response functions and their frequency domain.

response functions for a particular event can be either *asymmetric* (Figure 5, top), featuring only buildup or only decay, or *symmetric*, featuring both sides, not necessarily of the same shape (Figure 5, middle). Both types are important to recognize, since they allow to reason about events semantics.

4.1.2 Modeling Events Importance

We model events by using a piecewise linear approximation of event importance $e(t)$. This model introduces a set of meaningful parameters, that can succinctly describe event importance, namely, *buildup and decay rates*, *longitude* of an event and its *maximum importance* level. We show the most basic shapes of event importance in Figure 1 (bottom) and describe them below. Accordingly, there are rectangular (Event 2) and triangular (Events 1 and 3) models, which can also merge to trapezoid-like shapes (Event 3’).

The Rectangular model of $e(t)$ is suitable for long-duration events with roughly constant importance and coverage, like “Olympics”. This model is represented using a step function with the constant height of n_0 and the longitude τ_e : $e(t) = n_0, 0 \leq t \leq \tau_e$.

The Triangular model adds the parameters of buildup a and decay b , when $e(t)$ is of varying importance to mass media during its period. Accordingly, we can represent it using a piecewise linear function, which originates at time $t = 0$, and reverses its direction at $t = t_0$: $e(t) = at, 0 \leq t \leq t_0$; $e(t) = (a+b)t_0 - bt, t > t_0$.

We show the shapes of the resulting news functions in Figure 6. In this example, we used news events of three different longitudes ($\tau_e = 0.25, 0.5$ and 1.0 day) and selected the parameters of response functions so that $n(t)$ would reach the same amplitude after 0.5 days for an event with a constant importance (as can be seen in the second event). We observe that a varying time length of events in this model results in varying sharpness of news frequency peaks. Correspondingly, short-time events have a shorter period of actuation and long-time events eventually result in the saturation of news media. In all cases, the period of relaxation of news media is mainly characterized by the shape of the corresponding media response function.

According to our model, the height of the event on the event sequence $e(t)$ indicates its importance, while the length describes its longitude. Events can be of constant importance, like those shown in Figure 6 (left side), or of varying importance, shown in Figure 6 (right side). In both cases, the time longitude as well as the maximum importance can be different for different events even for the same topic.




Figure 6: Convolution of rectangular and triangular event models.

4.1.3 News Deconvolution

Deconvolution is the opposite process to convolution [3]. With this procedure we are able to recreate the original event importance sequence, as shown in Figure 1, bottom. Moreover, deconvolution helps to detect and separate nearby, but distinct events, which may otherwise be considered as a single event due to their overlap.

However, deconvolution can be expressed in the functional form only for a limited number of response functions (e.g. exponential, linear), while our framework is designed to support any kind of finite integrable functions. Therefore, we take an approach to deconvolution, which relies on the frequency domain of the signals, as explained below.

The Convolution theorem states, that the Fourier transformation of a time-domain convolution of two series is equal to the multiplication of their Fourier transformations in the frequency domain:

$$\mathcal{F}\{n(t)\} = \mathcal{F}\{e(t) * mrf(t)\} = \mathcal{F}\{e(t)\} \cdot \mathcal{F}\{mrf(t)\} \quad (10)$$

According to this equation, for every cyclic frequency ω , we can obtain Fourier coefficients $e(\omega)$ of the original series $e(t)$ by dividing Fourier coefficients of the observed series $n(t)$ by the corresponding coefficients of the response function $mrf(t)$. During this process, we can also remove global trends from the time series by lowering the values of low-frequency components. Finally, we obtain a time-domain representation of $e(t)$ by performing an inverse Fourier transformation:

$$e(t) = \mathcal{F}^{-1}\{e(\omega)\} = \mathcal{F}^{-1}\{n(\omega)/mrf(\omega)\} \quad (11)$$

In this paper, we use the analytical Fourier representations of the three example response functions considered in our framework, although it is possible to use numerically-computed Fourier coefficients of *any arbitrary* decay function (symmetric or asymmetric).

Our method is also able to distinguish between symmetric and asymmetric $mrf(t)$, and use the appropriate one for deconvolution. For instance, let's consider the differences between the two models for the “Solar Eclipse” event from Twitter, illustrated in Figure 7. In the left part of the figure, we observe that the main event (shown in red) has a duration of one day, and was also discussed the day after (when photos were published). The symmetric deconvolution was able to extract these parameters and correctly match (green color) the observed volume of publications (blue color). For the asymmetric deconvolution, shown on the right, if we consider only the decay part of $mrf(t)$, we end up with the incorrect estimation of event's time and duration and also less correct prediction of the volume (green color, right side).




Figure 7: Symmetric (expected) and asymmetric (unexpected) response functions applied to “Solar Eclipse” event from Twitter.

Deconvolution with small decay parameters is fail-safe and loss-less as long as event shapes are not approximated or event importance is not becoming negative. However, only a high level of deconvolution helps to recover the original event importance, and to represent its dynamics using the desired model. Since higher than necessary deconvolution may result in smaller events becoming outcast by a “shade” from their preceding larger neighbors, we need to apply the largest possible, but still adequate level of deconvolution, using the parameter estimation as described below.

4.1.4 Parameters Estimation

A conventional method for estimating parameters of response functions relies on observing event decay shapes under the assumption that events are singular in time (i.e., they can be represented by Dirac's delta function). Parameter estimation in this case can be done by regression parameter fitting, following the peak detection and extraction of their slopes. First, we normalize the ascending or descending slopes of time series and then analyze them using either linear, power-law or exponential regression. Next, we take the average of the extracted parameters across the peaks, weighting them by regression error and level of importance.

Nevertheless, regression fitting can only be achieved for a few known response functions. In order to learn the parameters of *any* response function, we propose to optimize them over the fitness of our model for the entire time series (of news volume for a specific topic). A direct measure of fitness can be computed by averaging fitting errors of piecewise linear regression of event importance. However, this measure only improves with more powerful deconvolution, as more and more events obtain spike shapes, thus defeating our purpose. Therefore, we optimize the parameters using residual deconvolution errors, as follows.

First, we perform a deconvolution of $n(t)$ using a chosen response function with an initial set of parameters. Then, we approximate the resulting time series using a piecewise linear function, obtaining triangular $e(t)$. Following that, we repeat the process of deconvolution in the opposite way: convolving the approximated series with the same response function. Finally, the output is compared to the original time series, and the error of fitting is calculated. If the error of fitting improves by flipping to asymmetric response function, this type becomes chosen. The whole process repeats until the optimal parameters are identified. It is also possible to determine the best response function type for every peak and process the corresponding time interval individually. This choice can be based on the best fitness, or on the preferences of the given media and topic. In this paper, we apply deconvolution to the whole time series, to compare different response functions and the stability of their parameters, leaving per-peak processing for the future.

The proposed method is effective for identifying parameters for any event models and response functions. It is also efficient thanks to a concave error surface, where for small decay times the error is dominated by the event's model approximation, and for large decay times the error is caused by a longer than necessary response. Therefore, we apply gradient descent for estimating parameters.

4.2 Detecting Sentiment Shifts

In this paper, we are interested in sentiment measures that are sensitive to particular kinds of sentiment changes and that can also be correlated with events. Nevertheless, not many studies propose a suitable measure for opinion shifts, which can be analyzed coherently with the news time series in order to extract correlations. The particular methods which can be adopted to our problem are *sentiment volume* [10] and *contradiction level* [15], discussed below.

4.2.1 Sentiment Extraction

We extract texts for a particular topic using the Apache Lucene index, by querying it for documents that contain a given topic’s keywords. Analyzed texts are first cleaned from markup, such as urls, user names, hashtags, and quotes in the case of tweets. Then, we use the SentiStrength algorithm [11, 10] to extract sentiment polarities of texts. Obtained sentiment labels S are scaled to be in the range $[-1;1]$, and later aggregated as real values.

4.2.2 Sentiment Volume

Thelwall et al. [10] evaluate how twitter sentiment and its volume are changing before and after news events. By analyzing the peaks in sentiment, they show that the volume of negative sentiment is increasing just before an event, while there is an increase of positive sentiment at the event’s peak intensity. Their results indicate that external events usually lead to changes in sentiment and, more importantly, increase the contradiction level. One more observation made by the same authors is that the changes in sentiment are particularly small, making it necessary to apply more sophisticated methods for shift detection and properly align these small fluctuations with the underlying news time series.

Sentiment Volume is defined as the amount, or the sum of sentiments of the same polarity, expressed within a specified time interval. It captures bursts of a particular opinion polarity, e.g., *positive*:

$$s(t) = \sum_{i=1}^n S_i^+(t), \text{ or } s(t) = |S_i^+(t)| \quad (12)$$

4.2.3 Contradiction Level

Another suitable measure for sentiment shifts is the sentiment contradiction level [14, 15], which is based on the first statistical moments of sentiment values. The intuition behind this measure is that when the average sentiment value μ_S is close to zero, while the sentiment variance σ_S^2 is high, then the polarization of sentiments is high, indicating contradiction and diversity. Combining μ_S and σ_S^2 in a single formula, the authors propose the following measure for contradictions:

$$s(t) = \frac{\vartheta \cdot \sigma_S^2}{\vartheta + \mu_S^2} W(n) \quad (13)$$

where n is the number of sentiments, $\vartheta \neq 0$ is the normalizing constant, and W is a weight function that takes into account the significance of sentiment statistics involved in the calculation [15]. A distinct property of this measure is that it detects both changes of sentiment polarity as well as temporary shifts of sentiments.

4.3 Correlating News and Sentiments

We observe that different sentiment and news measures require different correlation methods $\rho(s, n, \delta)$, which also consider a time lag between the time series, δ . In the case of continuous time series, which usually deviate around their average values, we can use the Pearson cross-correlation coefficient, which is defined as the normalized covariance of two time series:




Figure 8: Correlated news and sentiment series from Twitter for movie “Hangover”, where bursts are extracted using deconvolution.

$$\rho(s, n, \delta) = \frac{\text{Cov}[n(t), s(t + \delta)]}{\sigma_s \sigma_n} \quad (14)$$

However, Pearson correlation is intended to determine a linear dependency between variables, which is hardly observable for bursty time series such as unexpected events or sentiment contradiction. Such time series do not have a definite average level, around which the movement is happening. Instead, their values are outbursting from the minimum level at some points in time. This behavior requires a special correlation technique, which takes as input only the bursty points within a specified time interval. Let’s assume that sets of bursts for sentiment and news are denoted as S_t and N_t respectively. Following this, any kind of binary similarity measure can be applied, for example *cosine similarity* ρ_C or *Jaccard coefficient* ρ_J :

$$\rho_C(s, n, \delta) = \frac{|S_{t+\delta} \cap N_t|}{||S_{t+\delta}|| ||N_t||}, \quad \rho_J(s, n, \delta) = \frac{|S_{t+\delta} \cap N_t|}{|S_{t+\delta} \cup N_t|} \quad (15)$$

In the above equations, intersecting bursts are determined according to some proximity region ξ , and in addition to counting the number of overlapping bursts, we can apply burst weighting, for example based on their magnitude.

While we can determine bursts of sentiments using simple methods, events require more robust methods, as can be seen in Figure 8, where a group of two (out of three) events in the right part of the time series are located on the monotonic slope of $n(t)$, and are not detectable by thresholds, or by derivative tracking. Remarkably, deconvolution can recover these and other “hidden” events and result in a more accurate correlation. It is also crucial to consider the level of event importance, obtained after the deconvolution, rather than the original news volume, as can be seen in Figure 2, where iPad 1 and iPad 2 announcements had different importance to people (red) but almost the same volume (blue) due to a trend. Thus, we select $n(t)$ peaks by their level of importance, that yields the highest correlation.

We also consider that sentiment changes may be preceded or followed by news events with some time lag. In order to align the two sequences, we have to determine this time lag, which is different for different topics and media. Here again, deconvolution helps identifying correct event timings and, therefore, time lags. It shifts peaks backwards in time to the extent determined by a response function and event dynamics. The optimal time lag Δ can be determined by maximizing the cross-correlation coefficient over a range of positive (reaction) and negative (anticipation) time lags δ :

$$\Delta = \underset{\delta}{\text{argmax}} [\rho(s, n, \delta)] \quad (16)$$

An unconstrained optimization is inefficient, since it requires computing correlation for every parameter δ , whereas a heuristic optimization can produce wrong results, when news time sequence is a repetitive process. To overcome this, the absolute time lag must not exceed the first maximum of the news autocorrelation function.

5. EXPERIMENTAL EVALUATION

The main goal of our experimental evaluation is to study the properties of the real data and evaluate the proposed and existing models and their assumptions. We analyze the news dynamics and evaluate our models on several social media datasets with different characteristics. The agenda of our evaluation is the following:

- Compare accuracy to previous models;
- Evaluate the proposed response dynamics;
- Check the correlation of events and sentiments;

Our *Meme* dataset [6] consists of the top 100 meme time series² and approximately 500 peaks, that represent various events of endogenous nature, which circulated in the blogosphere in the period from August 2008 to May 2009. While some of the analyzed memes have a connection to real news events, the largest part of them are just sticky phrases gaining popularity from time to time. The analyzed time series have a granularity of aggregation of 4 hours and varying lengths (typically, one month).

For our *Twitter* dataset, we selected 30 topics³, which featured the most prominent events for the period of half a year, from June 2009 till December 2009. The dataset contained approximately 7 million tweets in total and over 400 peaks during the events. We extracted time series of volume and sentiments using a 1 day aggregation. Twitter data has very different properties when compared to the output of blogs or news media. First, this platform has a distinct bias towards current events and temporal activity of users. So usually there is no accumulated interest, since any activity fades out to zero after a short time. Second, there are different types of dynamics present at the same time: *daily activity* and *rending activity*. Whereas the first one is largely driven by work schedules in different time zones, the second one demonstrates a more clear pattern of event interest and is our main subject.

To compare the accuracy of our models for different media, we perform news volume deconvolution using automatically extracted parameters (as described in Section 4.1.4), and then quantify the accuracy of fitting in a way that is comparable across different time series (and across different peak heights). Accordingly, errors in fitness are measured as RMSE for every peak, and then normalized by its height, so that the results can be averaged and compared across peaks.

5.1 Evaluation of Meme Dataset

Taking into account our findings in Section 3.2, we used dynamics equation (1) to predict the values of every peak by applying it on the volume and time accumulated since the peak's start time. Considering the linear form for the imitation factor expressed as $f(x) = a + bx$, and $\delta(t) = t^{-\alpha}$ ($\alpha > 1$), we have:

$$\begin{aligned} \frac{dx}{a+bx} &= ct^{-\alpha} dt \\ \frac{1}{b} \ln |a+bx| &= \frac{c}{1-\alpha} t^{1-\alpha} + C \\ f(x) = a+bx &\sim \exp\left(\frac{bc}{1-\alpha} t^{1-\alpha}\right) \\ n(t) = \frac{dx}{dt} &\sim \exp\left(\frac{bc}{1-\alpha} t^{1-\alpha}\right) t^{-\alpha} \end{aligned} \quad (17)$$

In Figure 9 (a) we demonstrate an example prediction of this model in the case when $\alpha = 3.3$ and $bc = 3.3$. We observe that the first peak's buildup and decay can be matched using these parameters only approximately, while the shapes of all the subsequent peaks cannot be matched at all.

²The complete meme dataset is available at: <http://www.memetracker.org/>

³The list of selected topics is available upon request.

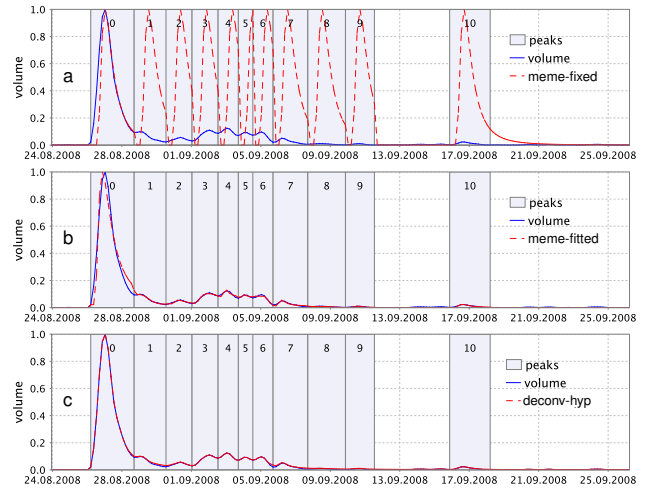


Figure 9: “No way, no how, no McCain!” meme from [6]. Peaks are approximated by meme model [6] using fixed (a) and fitted (b) parameters, and by deconvolution using fixed parameters (c).

Fitting equation (1) to individual peaks using *least squares* regression yields much better accuracy for small peaks (Figure 9 (b)), but does not match the sharpest peaks, where buildups and decays seem to require different model parameters. In the above experiment, fitting employed four parameters, which were independent among peaks and had different values for every peak. Two of these parameters are effectively stretching the model to accommodate real data: base level (a) indicates the volume accumulated before the peak, and scale (b) serves the purpose of fitting the peak's height, to counterweight the normalization effect caused by dividing on volume in equation (1). Normalizing the predicted volume over the top 100 time series (the concurrency factor) did not yield any improvement neither for fixed, nor for individually fitted peak parameters. This can either indicate that the time series evolve independently of each other, or that the top 100 time series are not sufficient to cover the whole publishing activity.

Overall, we observed an irregularity of meme model's parameters for different peaks of the same time series. In contrast, our deconvolution model uses a single set of parameters for the entire time series, which are automatically estimated by performing a peak slope regression (using a single or multiple peaks). Figure 9 (c) demonstrates the output of a hyperbolic deconvolution model, which fits the time series consistently across all peaks. The output in this case is the time series, constructed by performing an inverse process of convolution over the deconvolution-estimated time series of event importance.

5.1.1 Evaluation of Accuracy

First, we compare the observed error distributions between the meme and hyperbolic deconvolution models in Figure 10. A hyperbolic deconvolution model was chosen as the closest match to a meme model regarding the predicted shapes of peak slopes (refer to Equation 17). Moreover, our model used triangular $e(t)$ shapes, automatically extracted using a deconvolution with parameters $\tau = 0.8$ days and $\alpha = 2.0$, fixed for all time series. The deconvolution model reached an average RMSE of 0.11 with the standard deviation of 0.09. The meme model demonstrated an average RMSE of 0.22 with standard deviation of 0.09. Taking into account the artifacts of data aggregation, noise and deviation of peak starting times, the error level of 0.2 can still be considered as acceptable




Figure 10: Accuracy of meme (fit) and deconv (fixed) models.




Figure 11: Deconvolution models accuracy for Meme dataset.

for the meme model. Nevertheless, it fitted more than half of the evaluated peaks with larger errors than our model, which used a single set of parameters for *all* time series.

We note that choosing smaller decay times results in smaller differences between event importance (deconvolved volume) and volume. On the other hand, this introduces more errors to the estimation of $e(t)$ shapes (performed by piecewise linear regression) and therefore results in larger errors for the output approximation. Choosing larger decay times, can also induce approximation errors due to the omission of smaller peaks from $e(t)$. Nevertheless, the output of our model is stable with respect to parameter variation (as demonstrated by our experiment in Figure 10, where fixed parameters still yield a good approximation quality). Thus, it is possible to perform deconvolution using the estimated parameters achieving nearly the same performance as when using the optimal parameters.

The performance of deconvolution models on meme dataset is depicted in Figure 11. We observe that the linear and hyperbolic deconvolution models reached the best accuracies with average errors of 0.03 and 0.05 respectively, and standard deviations of 0.06 (thus, the difference in performance is not significant). However, while the linear model had an average parameter τ of 0.5 (very small), the average parameters τ and α of the hyperbolic model were 0.36 and 2.8 respectively. A more careful evaluation of the estimated parameters reveals that the hyperbolic model has almost the same parameter τ of 0.36 ± 0.03 across all topics in the dataset, but quite different α of 2.8 ± 0.96 . A similar pattern is observed for the triangular model, where a single decay parameter τ has values in the range 0.5 ± 0.3 . The exponential model demonstrated an average RMSE of 0.11 with standard deviation of 0.09, almost the same to our previous experiment with fixed parameters, but in this case the estimated parameter τ was (on average) equal to 0.65, thus trying to approximate a usually much steeper hyperbolic response. From these observations we conclude that meme data is more likely to have a hyperbolic response pattern. But while the dominance of linear and hyperbolic response functions is clearly visible, their parameters are significantly different from peak to peak.

5.2 Evaluation of Twitter Dataset

We analyzed 30 time series from Twitter using $1d$ aggregations. The results of our evaluation, presented in Figure 12, demonstrate a good performance for the linear model, though the average parameter τ was 2.5 with a deviation of 0.9, indicating that a rather high fraction of peaks had small τ . Since in our dataset we usually have a small number of samples per peak (8-10), it is not possible to fully verify the hypothesis of the linear model. The obtained results only prove that most events have a fast rise and drop of importance around their peaks. The two other models (hyperbolic and expo-




Figure 12: Deconvolution models accuracy for Twitter dataset.

ponential) have more powerful estimated decay parameters than the linear model but also larger errors ($\alpha = 2.4 \pm 0.5$, $\tau_h = 1.3 \pm 0.6$ and $\tau_e = 1.5 \pm 0.7$). Finally, we observe that both models encountered large variations of the error, probably indicating the existence of *different* kinds of response dynamics for topics or different response parameters for various events during the same time series.

5.3 News and Sentiments Correlation

In this section we evaluate correlation between sentiment shifts and events, where events are extracted with the help of deconvolution, and sentiment shifts are determined using the measures of positive and negative sentiment volumes $s^+(t)$ and $s^-(t)$, and sentiment contradiction level $C(t)$.

To evaluate the correlations, we picked 22 out of 30 selected topics in our Twitter dataset, which contained a sufficient number of events, and were suitable for computing bursts correlation. The results of our evaluation are summarized in Table 2, where for each time series and sentiment measure we present the values of Cosine similarity ρ_C and time lag Δ of the best correlation. Additionally, we report the best-fit response dynamics "Type", and the average number of overlapping bursts "Num". We do not report Jackard coefficient ρ_J , since for highly-overlapping sets it shows the same behavior as ρ_C .

We observe that negative sentiments demonstrate the highest overall correlation, followed by contradictions and positive sentiments. Taking a closer look at the data, we see that positive sentiments are usually preceding the expected events (correlated with a negative lag), whereas negative sentiments are observed after or during these events (positive lag). For instance, such behavior is observed in our table for events surrounding the movie sequels "Hangover", "Ice Age" and "Harry Potter", which had positive expectations before (during) the event, and negative sentiment afterwards. Anticipated events like "LCROSS", "Super Bowl" or "Leica" also have positive expectations, but in this time connected to events. Another group of anticipated events, like "CERN LHC", "Beer Summit", "Aquila Summit" and "Transformers" features negative expectations and positive outcomes. Interestingly, heavily promoted events like new camera announcement for "Leica" on 09.09.2009, or availability of "Google Wave", have initially positive sentiments, but then receive a mixture of controversial opinions during their peaking intervals. In addition, some of unexpected events, like "Iran Election", "Iraq", "Obama", "Gmail Blackout" seen a bump of negative sentiments after their appearance, while Michael Jackson death flipped initially controversial (negative and positive) sentiments to positive.

Overall, we observe that the particular types of sentiment changes are connected to the particular behavior of events, and the same is true for their news dynamics. While different topics and media can be better described with different response functions, unusual events can still occur in the course of time, affecting sentiments in their own way. Therefore, it makes sense to organize the dependency modeling component to operate at the level of events, rather than topics. This component still needs to monitor correlations for each topic, but this time filtering the events according to their type and particular characteristics, before trying to detect the relevant sentiment shifts in their proximity.

Topic Name	Type	Num	Contradiction		PosVolume		NegVolume	
			ρ_C	lag	ρ_C	lag	ρ_C	lag
Measures								
Hangover	lin	20	0.68	0.0	0.65	0.0	0.60	2.0
Ice Age	exp	22	0.58	0.6	0.66	-1.0	0.65	1.6
Harry Potter	exp	13	0.42	1.7	0.57	0.0	0.45	1.7
LCROSS	hyp	12	0.65	-0.9	0.73	-1.5	0.70	-1.4
CERN LHC	hyp	21	0.68	1.0	0.61	1.0	0.54	-1.4
NASA	hyp	6	0.67	-1.5	0.72	1.5	0.74	-1.5
Michael Jackson	hyp	14	0.54	-1.3	0.49	1.5	0.51	-1.1
Eclipse	hyp	6	0.67	-1.4	0.42	2.1	0.50	1.9
Swine Flu	exp	19	0.58	0.6	0.59	0.0	0.56	-0.5
Barack Obama	hyp	17	0.60	-1.6	0.55	0.0	0.47	0.6
Iran Election	hyp	17	0.51	-0.5	0.47	0.0	0.53	1.6
Iraq	hyp	22	0.61	-1.6	0.58	-1.6	0.74	0.6
Follow Friday	lin	20	0.83	-1.3	0.62	0.7	0.75	-1.4
"Aquila Summit"	exp	30	0.79	0.5	0.83	0.9	0.96	-0.5
"Beer Summit"	hyp	6	0.71	0.8	0.62	0.9	0.71	-0.8
Transformers	exp	16	0.56	-1.8	0.58	0.6	0.51	-1.9
Facebook	hyp	15	0.60	-1.2	0.50	0.8	0.56	-1.3
TwitterPeek	hyp	6	0.72	-1.3	0.72	0.0	0.77	0.0
Leica	hyp	22	0.70	0.0	0.60	-1.7	0.80	-0.5
Gmail	hyp	20	0.71	0.6	0.55	-1.9	0.67	0.6
Super Bowl	lin	20	0.71	0.0	0.53	-1.6	0.71	0.0
Google Wave	hyp	11	0.50	0.5	0.67	-1.0	0.51	0.5

Table 2: Correlation between sentiment shifts and events in Twitter.

6. CONCLUSIONS

Our evaluation of news dynamics and their impact on sentiments is the first systematic work in this direction, which employs a thorough and principled modeling of news distribution in various media and studies news interaction with sentiments. We develop a novel model of news event dynamics, which allows to capture meaningful and important characteristics of news events. Our model can accommodate various response functions, suitable for different cases, which should not necessarily be expressed as a differential equation, but can be learned from the data. The results obtained by applying our methods to different real datasets confirm their robustness and universality. Nevertheless, they also reveal that we need to address several more major challenges.

First, we observe that while different media have preferences for particular response dynamics, these are often determined by event types and topics. Thus, we need to extend our method of news volume deconvolution so that it will automatically determine the best model for every particular event and process the corresponding time interval individually. The choice of the suitable deconvolution model can be based on the (learned) preferences of the given media and topic, but should also take model fitness into account. Above all, this involves a refinement of the events importance model and development of a robust and precise deconvolution optimization strategy.

We also observe the existence of different parameters of response dynamics for various events even during the same topic time series. Each of these events has a different importance, and also a different impact on sentiments, suggesting that it is possible to predict sentiment changes. In order to do that, we need to take into account the type of response dynamics (*mrf*) in addition to the event's importance level, creating a more elaborate causality model. Still, predicting the types of these possible changes requires building a database of event and sentiment shift profiles, and constructing a classifier model based on these features. We aim at creating such a dataset as our next step towards the final solution.

Finally, we speculate that sentiment shifts may be caused by events on related topics, in addition to events on the same topic. This leads to the necessity for a more advanced correlation and causality modeling, in order to predict sentiment shifts across related topics.

7. REFERENCES

- [1] S. Asur, B. A. Huberman, G. Szabó, and C. Wang. Trends in social media: Persistence and decay. In *ICWSM*, 2011.
- [2] R. Crane and D. Sornette. Robust dynamic classes revealed by measuring the response function of a social system. *Proceedings of the National Academy of Sciences*, 105(41):15649–15653, 2008.
- [3] K. Gaikovich. *Inverse problems in physical diagnostics*. Nova Science Publishers, 2004.
- [4] X. Hu, L. Tang, J. Tang, and H. Liu. Exploiting social relations for sentiment analysis in microblogging. In *WSDM*, pages 537–546, New York, NY, USA, 2013. ACM.
- [5] J. Lehmann, B. Gonçalves, J. J. Ramasco, and C. Cattuto. Dynamical classes of collective attention in twitter. In *WWW*, pages 251–260, 2012.
- [6] J. Leskovec, L. Backstrom, and J. Kleinberg. Meme-tracking and the dynamics of the news cycle. In *KDD*, pages 497–506, New York, NY, USA, 2009. ACM.
- [7] Y.-R. Lin, D. Margolin, B. Keegan, and D. Lazer. Voices of victory: A computational focus group framework for tracking opinion shift in real time. In *WWW PE Track*, 2013.
- [8] D. Sornette, F. Deschâtres, T. Gilbert, and Y. Ageon. Endogenous versus exogenous shocks in complex networks: An empirical test using book sale rankings. *Phys. Rev. Lett.*, 93:228701, Nov 2004.
- [9] I. Subasic and C. Castillo. Investigating query bursts in a web search engine. *Web Intelligence and Agent Systems*, 11(2):107–124, 2013.
- [10] M. Thelwall, K. Buckley, and G. Paltoglou. Sentiment in twitter events. *JASIST*, 62(2):406–418, 2011.
- [11] M. Thelwall, K. Buckley, G. Paltoglou, D. Cai, and A. Kappas. Sentiment in short strength detection informal text. *JASIST*, 61(12):2544–2558, 2010.
- [12] M. Tsytarau, S. Amer-Yahia, and T. Palpanas. Efficient sentiment correlation for large-scale demographics. In *SIGMOD*, pages 253–264, 2013.
- [13] M. Tsytarau and T. Palpanas. Survey on mining subjective data on the web. *Data Mining and Knowledge Discovery, Special Issue on 10 Years of Mining the Web*, pages 1–37, 2011. ISSN 1384-5810.
- [14] M. Tsytarau, T. Palpanas, and K. Denecke. Scalable discovery of contradictions on the web. In *WWW*, 2010.
- [15] M. Tsytarau, T. Palpanas, and K. Denecke. Scalable detection of sentiment-based contradictions. In *DiversiWeb Workshop, WWW 2011, Hyderabad, India*, 2011.
- [16] A. Tumasjan, T. O. Sprenger, P. G. Sandner, and I. M. Welp. Predicting elections with twitter: What 140 characters reveal about political sentiment. In W. W. Cohen and S. Gosling, editors, *ICWSM*. The AAAI Press, 2010.
- [17] F. Wu and B. A. Huberman. Novelty and collective attention. *Proceedings of the National Academy of Sciences*, 104(45):17599–17601, 2007.