



Research Center
Finance & Information Management



Project Group
Business & Information
Systems Engineering

Detecting Online Firestorms in Social Media

by

Benedict Drasch¹, Johannes Huber, Sven Panz¹, Florian Probst

in: Proceedings of the 36th International Conference on Information Systems
(ICIS), Fort Worth, USA, December 2015

¹ Students of the Elite Graduate Program “Finance- & Information Management”,
University of Augsburg, Germany

WI-455

University of Augsburg, D-86135 Augsburg
Visitors: Universitätsstr. 12, 86159 Augsburg
Phone: +49 821 598-4801 (Fax: -4899)

University of Bayreuth, D-95440 Bayreuth
Visitors: Wittelsbacherring 10, 95444 Bayreuth
Phone: +49 921 55-4710 (Fax: -844710)



Detecting Online Firestorms in Social Media

Completed Research Paper

Benedict Drasch

FIM Research Center,
University of Bayreuth
95440 Bayreuth, Germany
benedict.drasch@fim-rc.de

Johannes Huber

FIM Research Center,
University of Augsburg
86135 Augsburg, Germany
johannes.huber@fim-rc.de

Sven Panz

FIM Research Center,
University of Augsburg
86135 Augsburg, Germany
sven.panz@fim-rc.de

Florian Probst

FIM Research Center,
University of Augsburg
86135 Augsburg, Germany
florian.probst@fim-rc.de

Abstract

As social media has increased the reach and speed of electronic word-of-mouth (eWOM), so it has intensified customers' exposure to negative eWOM. Consequently, companies increasingly suffer from massive outbursts of negative eWOM, known as online firestorms. Because of their dynamics, it is nearly impossible to stop online firestorms if their emergence is not detected promptly. However, well-founded approaches that provide automated, real-time detection are missing. We design an Online Firestorm Detector that includes an algorithm inspired by epidemiological surveillance systems. Real-world data from a firestorm suffered by Coca-Cola is used to prove the utility and validity of the proposed approach. We show that online firestorms can be reliably detected shortly after the first piece of related negative eWOM has been generated, and that the number of false alarms is low. A comparison with competing artifacts shows that the Online Firestorm Detector is superior to approaches that could be alternatively used.

Keywords: Word-of-mouth, social media, online firestorm, information diffusion, design science

Introduction

When prospective customers search for information about products and services, publicly available online reviews, ratings, and critiques of fellow consumers are increasingly important (Chen and Xie 2008; Dellarocas 2003; Dellarocas et al. 2007; Moon et al. 2010). As more than 85% of online purchases are currently driven by such ratings and reviews (Stratus Contact Solutions 2014), this so-called electronic word-of-mouth (eWOM) has made social media a particularly important factor in companies' marketing communications (Albuquerque et al. 2012; Faase et al. 2011; Forman et al. 2008). Hence, companies support the creation of customer-to-customer interactions in social media channels such as Facebook and Twitter (Harris and Dennis 2011; Poynter 2008). Indeed, prior research has shown that positive eWOM generated in social media creates tremendous business value (cf. e.g., Goh et al. 2013; Moe and Trusov 2011; Rishika et al. 2013).

However, as social media has massively increased both the reach and speed of eWOM diffusion, so it has accelerated and intensified customers' exposure to negative eWOM. An early famous example of customers being exposed to massive negative eWOM is often referred to as "Dell Hell" (Jarvis 2005). On June 21st 2005, a blogger described his disappointing experience with Dell's in-home service. Within hours, an increasing number of people had read the post, agreed with the blogger, and shared their own negative experiences with Dell. This domino effect has been cited as a major factor leading to a steep drop in Dell's customer satisfaction rating, and even its share price (Furfie 2008). This example illustrates the considerable risks of eWOM: If the sentiment turns negative, it can spread like a firestorm, leading to a considerable loss of customers and damage to the firm's reputation (Hennig-Thurau et al. 2010; Mochalova and Nanopoulos 2014; Pfeffer et al. 2014). Hence, practitioners and researchers coined the term "online firestorm", which can be defined as "[...] the sudden discharge of messages containing negative [e]WOM and complaint behavior against a person, company, or group in social media networks" (Pfeffer et al. 2014, p. 118). Over recent years, companies from diverse industries have suffered from such online firestorms and their negative consequences (cf. Pfeffer et al. 2014).

With the ability to reach thousands of (potential) customers in a very short time, it is nearly impossible to stop the diffusion of negative eWOM if the emergence of an online firestorm is detected too late (Stich et al. 2014). Hence, a timely intervention is crucial to avoid (in the best case) the actual outburst of an online firestorm, or at least to initiate countermeasures as soon as possible (e.g., by showing public regret and apologizing, cf. Munzel et al. 2012). Because of the rapid nature and huge volume of eWOM, this can only be achieved by automated, real-time detection approaches. However, to the best of our knowledge, there is no well-founded approach for the automated, real-time detection of online firestorms in social media. Related prior research on sentiment analysis has confirmed the negative economic consequences of negative eWOM, but did not focus on the detection of online firestorms, whilst studies in the field of information diffusion in social networks focused on simulating the spread of negative eWOM on a user-to-user path basis. However, the question of *when* to trigger an alarm if negative eWOM spreads across a network has not yet been addressed.

We therefore design an IS artifact, that is, an Online Firestorm Detector based on an algorithm inspired by well-established research on epidemiological surveillance systems. The Online Firestorm Detector's validity and utility are empirically evaluated using real-world data from an online firestorm suffered by Coca-Cola on Facebook. Additionally, we substantiate the quality of our Online Firestorm Detector by comparing its performance to competing artifacts that could be alternatively used. Thus, we account for changes in online customer-to-customer interactions that have been enabled by information systems (IS) (cf. e.g., Libai et al. 2010) and their potential negative consequences such as online firestorms. This leads to theoretical and practical contributions: From a theoretical perspective, we enrich existing IS and marketing literature on the analysis of eWOM in social media to avert its potential dark side, demonstrate that research from the field of epidemiology can serve as a valid theoretical basis in the context of eWOM diffusion in social media, and extend the understanding of online firestorms by showing that negative eWOM is not the only factor that should be considered when detecting online firestorms. Thus, we contribute to the growing body of IS literature on social media, which has become an intensively researched topic over the last 10 years (cf. e.g., Berger et al. 2014; Heidemann et al. 2012; Probst et al. 2013). From a practical perspective, we show that common lightweight solutions are unsuitable for the reliable detection of online firestorms in social media, and provide a ready-to-use artifact that enables companies to mitigate risks from social media engagements.

This paper is based on the Design Science Research paradigm, and in particular on the publication schema of Gregor and Hevner (2013), which draws on the guidelines of Hevner et al. (2004) and the process model of Peffers et al. (2007). We perform five main steps: (1) We have discussed the “*purpose and scope*” of our artifact and its “*relevance*” to business practice within this introduction. (2) In the next section, we specify the problem context in more detail, and present findings from prior empirical research on the economic effects of positive and negative eWOM, which further emphasize that the detection of online firestorms constitutes an “*important and relevant business problem*”. Moreover, we focus on relevant existing “*descriptive and prescriptive knowledge*” by discussing related work on the diffusion of eWOM, epidemiological surveillance, and anomaly detection, which informs our design process. (3) We then present our “*design artifact*”, referred to as the “Online Firestorm Detector”. (4) Subsequently, we provide a “*rigorous design evaluation*” to demonstrate the validity and utility of the artifact using real-world data. The quality of our Online Firestorm Detector is evidenced by a comparison with “*competing artifacts*” that could be alternatively used. Afterwards, we critically “*discuss and reflect*” on our artifact in terms of its contribution, limitations, and future research. (5) Finally, we conclude with a brief summary.

Problem Context and Related Work

In this section, we first define the problem context of online firestorms in social media. Second, we provide an overview of existing findings from relevant related work within the contexts of information diffusion, epidemiological surveillance, and anomaly detection.

Problem Context: Online Firestorms in Social Media

Today, social media are the predominant platforms for communication and interaction between companies and customers, as well as among customers themselves (Goh et al. 2013; Kietzmann et al. 2011). Prior literature has classified social media into six different categories: (Micro-)blogs, online social networks (often also called social networking sites), virtual social worlds, collaborative products, content communities, and virtual game worlds (Kaplan and Haenlein 2010). Taking their common characteristics into account, social media can be defined as a “[...] group of Internet-based applications that build on the ideological and technological foundations of Web 2.0, and that allow the creation and exchange of User Generated Content” (Kaplan and Haenlein 2010, p. 61). A particularly important part of User Generated Content is eWOM, which can be defined as “[...] any positive or negative statement made by potential, actual, or former customers about a product or company, which is made available to a multitude of people and institutions via the Internet” (Hennig-Thurau et al. 2004, p. 39). In social media, the underlying network is based on technical features that allow users to build online relationships with many other users (e.g., Facebook friends) and communicate intensively among one another (e.g., via wallposts and comments in Facebook). As a result, users form dense network clusters (Benevenuto et al. 2009; Mislove et al. 2007; Wilson et al. 2009). Within these clusters, the information flow is usually relatively constant and unrestrained. There is typically a very short period of time until the next piece of information (Pfeffer et al. 2014), which supports the fast spread of information and provides the fuel for online firestorms (Lotan 2012). Consequently, an enormous number of people can be reached by eWOM within a short period of time (Pfeffer et al. 2014).

Characteristics of the Diffused Information

The diffusion of eWOM in social media depends on the characteristics of the information itself (Bampo et al. 2008; Goh et al. 2013), such as the type of emotion, sentiment, or level of physiological arousal (Berger and Milkman 2012). Negative, traditional word-of-mouth (WOM) has been found to be published more often and to influence users stronger than positive WOM (Anderson 1998), because individuals consider negative WOM to be more important (Rozin and Royzman 2001). Prior research has shown that individuals perceive also eWOM containing negative emotions such as awe, anger, or anxiety as more important and credible than positive eWOM (Berger and Milkman 2012). Therefore, eWOM with negative sentiment is particularly likely to diffuse more widely and with more extreme consequences than positive eWOM. Companies thus stand to suffer serious harm from the outburst of a massive amount of negative eWOM, that is, an online firestorm. As Table 1 shows, the existing literature broadly confirms the negative effects of eWOM with negative sentiment on economic measures such as revenue and customer cash flows. Tirunillai and Tellis (2012) even found that negative eWOM could decrease companies’ stock values.

Thus, negative eWOM can become a severe problem, if the overall sentiment toward a company turns negative. However, whereas research on the sentiment of eWOM in social media has confirmed the importance of avoiding the massive spread of negative eWOM, approaches that enable the automated, real-time detection of online firestorms have not been investigated in this line of research.

Authors	Context	Economic Measure	Pos. eWOM	Neg. eWOM
Chen et al. (2004)	Books	Sales rank	↑	↓
Chevalier and Mayzlin (2006)	Books	Sales rank	↑	↓
Chintagunta et al. (2010)	Movies	Revenue	↑	↓
Dhar and Chang (2009)	Music	Sales rank	↑	↓
Duan et al. (2008)	Movies	Revenue	↑	↓
Liu (2006)	Movies	Revenue	↑	↓
Ludwig et al. (2013)	Books	Conv. rate	↑	↓
Moe and Trusov (2011)	Beauty products	Revenue	↑	↓
Sonnier et al. (2011)	Tech. products	Revenue	↑	↓
Goh et al. (2013)	Apparel retailer	Revenue	↑	↓
Tirunillai and Tellis (2012)	Tech. products	Stock return	✱	↓

↑ positive, ↓ negative, ✱ no influence on dependent variable confirmed

Information Diffusion within Social Media and its Similarity to Epidemics

Research on information diffusion within social networks is closely related to the study of online firestorms. Prior research on the diffusion of eWOM in communication networks studied the so-called “influence maximization problem” (cf. Domingos and Richardson 2001). This refers to the combinatorial optimization problem of identifying the target set of influential people that maximizes the information cascade in the context of viral marketing (cf. also Richardson and Domingos 2002). Based on these studies, Kempe et al. (2003, p. 138) investigated two of the “[...] most basic and widely-studied diffusion models” from research on the spread of diseases and viruses (epidemiology), that is, the linear threshold (LT) and the independent cascade (IC) model (cf. also Chen et al. 2009; Leskovec et al. 2007). In epidemiology, similar to information diffusion in social media, both the characteristics of the disease, such as transmission probability or complexity (Newman 2002), and “[t]he structures of human contact networks undoubtedly play a crucial role in [their transmission]” (Bansal et al. 2007, p. 881). In line with previous research, the spread of negative eWOM can be compared to the spreading behavior of pathogenic organisms (cf. e.g., Budak et al. 2011; Kempe et al. 2003): Negative eWOM (like pathogenic organisms) may be disseminated via (virtual) social links between social media users (i.e., humans). The infection rate is thereby highly dependent on the characteristics of the diffused information (i.e., emotion such as negativity, the complexity of transmission), the underlying network structure (i.e., population density and interconnectedness), and the process of attaining information (i.e., the means of transmission through the network). Hence, LT and IC models have been intensively applied to model information diffusion in the context of social media (for an overview cf. Probst et al. 2013). Thereby, the main focus has been to identify an initial set of influential users that maximize the spread of usually positive eWOM (cf. Probst et al. 2013). A few studies have extended these models to analyze the emergence and propagation of negative opinions, as “[t]he problem of limiting the effect of misinformation in a social network can be seen as similar to the problem of epidemics [...]” (Budak et al. 2011, p. 666). For instance, Chen et al. (2011) extended the LT model to simulate the natural behavior of people turning negative. Another study used the LT model to block the spread of a rumor in a network by selecting the right nodes and “injecting” the opposing opinion (He et al. 2012). Related work employed an extended IC model to identify users who are not surrounded by negatively influenced users in order to stop or restrict the spread of an online firestorm (Mochalova and Nanopoulos 2013; Stich et al. 2014). Taken together, existing research shows that the diffusion of negative eWOM in social media can be simulated and

influenced on a user-to-user path basis. However, the simulation of diffusion paths is not sufficient to trigger an alarm if some critical threshold of negative eWOM in an entire network is exceeded. Hence, these models do not allow for the automated, real-time detection of online firestorms.


Anomaly Detection

Whereas research on the diffusion of information within social media focuses on user-to-user paths, the field of anomaly detection provides approaches that allow the criticality of information diffusion to be investigated from a global perspective. In general, early detection is based on signal detection theory (cf. e.g., Kay 1998) and decision theory (cf. e.g., Von Neumann and Morgenstern 1947), whereas anomaly detection approaches address the specific “[...] problem of finding patterns in data that do not conform to expected behavior” (Chandola et al. 2009, p. 1). Classification-based anomaly detection techniques (e.g., Bayesian networks, rule-based systems, nearest neighbor-based techniques) operate under the assumption that a classifier is able to “[...] distinguish between normal and anomalous [instances]” (Chandola et al. 2009, p. 19). For example, Bayesian networks estimate the posterior probability of an instance belonging to a normal or anomalous class (e.g., Wong and Cooper 2003). Rule-based systems determine *a priori* rules for understanding normal behavior and classifying instances that are not covered by these rules as anomalous (e.g., Agarwal 2005). Anomaly detection techniques using nearest-neighbor analysis classify instances according to their similarity or distance to classified instances (e.g., Lin et al. 2005). However, these existing anomaly detection techniques cannot be easily transferred to the detection of online firestorms in social media. First, these techniques are highly contextual in nature (Chandola et al. 2009), that is, approaches that work well in one context may perform poorly in a context such as eWOM in social media. Second, the detection of online firestorms requires the analysis of time-series data, which raises several statistical challenges such as seasonal and trend components that are not accounted for in the abovementioned approaches.

In contrast to classification-based techniques, parametric statistical anomaly detection (cf. e.g., Horn et al. 2001; Laurikkala et al. 2000) assumes that anomalous observations are not generated by the same stochastic model as normal observations (Anscombe and Guttman 1960; Brown 1971). Parameters for the underlying parametric model are estimated from normal data instances, and anomalous instances are classified, for instance, by applying hypothesis tests or evaluating instances against some threshold (Chandola et al. 2009). Other common procedures for detecting contextual anomalies (such as negative eWOM) also build on statistical, regression model-based techniques (cf. Chandola et al. 2009), which have been extensively investigated for time-series data (e.g., Fox 1972). The basic idea consists of two steps (cf. Chandola et al. 2009): First, a regression model is fitted; second, the residual for each test instance is used to determine the so-called anomaly score (cf. Anscombe and Guttman 1960). Based on this general principle, Farrington et al. (1996) described the “[...] standard reference method, routinely used [...]” (Guillou et al. 2014, p. 5026) that allows viral outbreaks of epidemics to be detected (cf. e.g., Freeman et al. 2013; Guillou et al. 2014). In contrast to other epidemiological surveillance approaches, Farrington et al. (1996) developed a general outbreak detection system. That is, without loss of general validity, the algorithm reliably performs on “[...] a very diverse range of organisms with different frequencies, trends and seasonality and [...] can be completely automated” (Noufaily et al. 2013, p. 1206). This is due to the fact that it focuses on general outbreaks, neglecting additional complexity (e.g., special, clustering). By virtue of its reliability and generalizability, the Farrington algorithm is actively used in several countries as part of their public health surveillance programs (Freeman et al. 2013; Hulth et al. 2010). Recently, Noufaily et al. (2013) improved the original algorithm by extending the period of the input data used in the regression analysis and adapting the re-weighting of unusual observations in the historical data used to forecast the expected development. These modifications brought about a substantial reduction in the false positive rate (i.e., type I errors) (Freeman et al. 2013). Hence, we base the design of our novel algorithm for the detection of emerging online firestorms on techniques for the early detection of outbreaks of infectious diseases.

Design of the Online Firestorm Detector

The proposed Online Firestorm Detector comprises three successive steps: (A) Monitoring social media and collecting eWOM, (B) conducting a sentiment analysis, and (C) detecting the emergence of online firestorms (cf. Figure 1). In step (A), the Online Firestorm Detector must perform the real-time monitoring and collection of company-related eWOM generated across social media. For most social media, publicly available application programming interfaces (APIs) can be utilized (cf. prototype in the following section). In step (B), the Online Firestorm Detector analyzes the sentiments of the collected eWOM. This is effectively a “[...] computational study of people’s opinions, appraisals, attitudes, and emotions toward entities, individuals, issues, events, topics and their attributes” (Liu and Zhang 2012, p. 415). Hence, the polarity of an opinion expressed in a given text unit (e.g., wallpost/comment on Facebook) can be identified (Thet et al. 2010), and the sentiment can be classified into categories such as “positive,” “negative,” and “neutral” (Kennedy 2012). In step (C), the actual emergence of an online firestorm must be detected. Whereas for steps (A) and (B) existing techniques can be leveraged, the detailed design of the novel algorithm utilized in step (C) is presented in the following.



As outlined in the previous section, there are no existing approaches for the automated, real-time detection of the emergence of online firestorms. However, related problems have been researched in the field of epidemiological surveillance systems (cf. previous section). As emphasized by Gregor and Hevner (2013, p. 347), it is common in IS research that “[...] effective artifacts may exist in related problem areas that may be adapted or, more accurately, exapted to the new problem context”. Therefore, our novel algorithm for the detection of emerging online firestorms builds on prior work on the early detection of outbreaks of infectious diseases by Farrington et al. (1996), which was further improved by Noufaily et al. (2013). We selected these specific algorithms for three reasons. First, the generalizability of the Farrington algorithm (cf. previous section) is an important factor in the context of online firestorms, as they are highly context-sensitive (e.g., depending on the company, topic, and so on). Second, the Farrington algorithm is highly reliable, and has been practically validated in many different contexts (e.g., different countries, different viruses). As a single online firestorm can significantly reduce a company’s value (cf. discussion in the previous section), this reliability is of utmost importance in the context of online firestorms. Third, the improvements made by Noufaily et al. (2013) have ensured that very few false alarms (type I errors) occur. This is particularly important, as “[t]he optimal level of sensitivity relative to specificity depends on the consequences of false alarms and the benefits of true alarms. These consequences are not fundamental properties of the detection method itself, but are specific to the use to which the detection method is being applied” (Wagner et al. 2001, p. 52). In the case of online firestorms, this means that a small delay (of the order of a couple of seconds) in detection time might be acceptable, whereas a high number of false alarms would significantly decrease confidence in the firestorm detection and cause alarms to receive less attention over time. Moreover, each time a false alarm was raised, unnecessary resources would be used to check whether it was in fact correct. The general algorithm behind the Online Firestorm Detector is depicted in Figure 2.

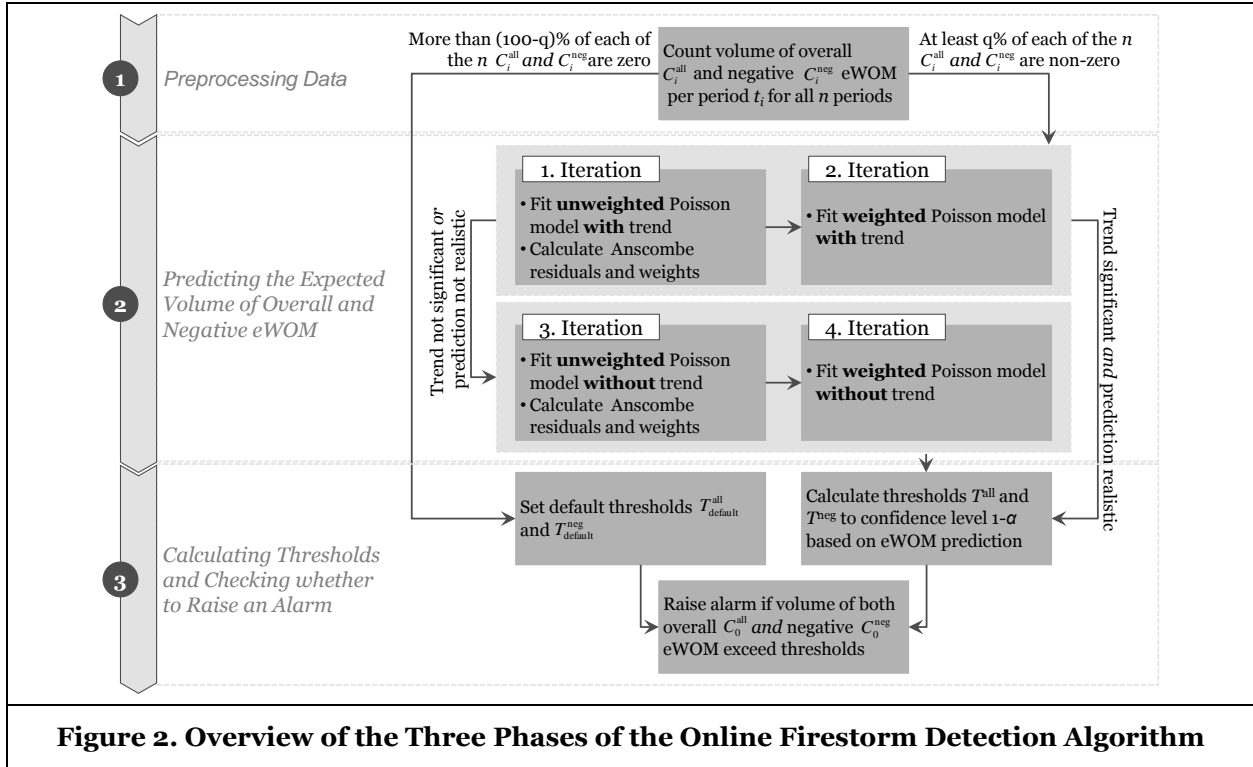
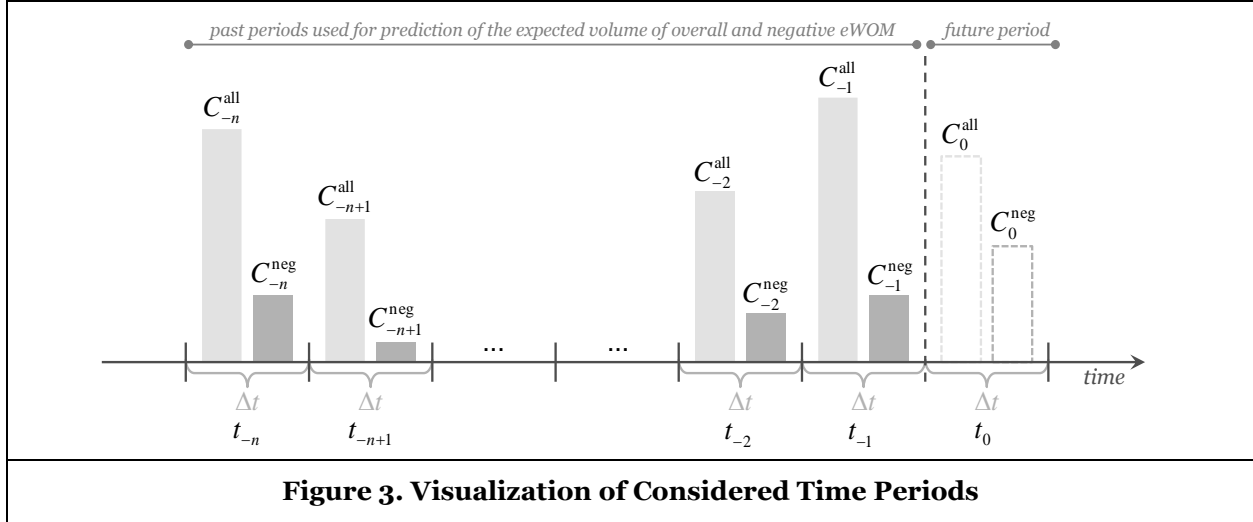


Figure 2. Overview of the Three Phases of the Online Firestorm Detection Algorithm

As shown in Figure 2, step (C) of the Online Firestorm Detector consists of three phases. In the first, the detector preprocesses data and checks whether a sophisticated prediction based on Poisson regression models can be applied, or whether the volume of eWOM in the periods under investigation has been too low. If a prediction can be made, the second phase of the algorithm identifies the best-fitting Poisson model. For this, a defined number of past periods are taken into account, and the algorithm runs several (un-)weighted Poisson regressions with and without a trend component. In the third phase, the best-fitting model is used to predict the expected volume of negative and overall eWOM for a defined future period and to calculate thresholds (to a predefined confidence level). If the actual volume of overall and negative eWOM exceeds these thresholds at any point in time within this future period, an alert is triggered. The detailed design of all three phases is described in the following.

Phase 1: Preprocessing Data

To analyze whether an online firestorm is emerging in the current period t_0 , we must first quantify the number of messages (i.e., volume of negative eWOM) in each past (micro-)period t_i , where $i = -n, -n + 1, \dots, -2, -1$, of length Δt . To do so, the algorithm counts the previously observed, collected, and analyzed (discrete) number of negative eWOM $C_i^{\text{neg}} \in \mathbb{N}_0 \forall i$. We consider the overall volume of eWOM $C_i^{\text{all}} \in \mathbb{N}_0 \forall i$ as well as the volume of negative eWOM. This is justified as follows: First, as the proposed approach aims to detect the outburst of online firestorms, we intuitively include negative eWOM to capture the cause of the online firestorm. Second, we incorporate a measure for the reach of negative eWOM, as the overall activity (total eWOM) on a company's fan page – whether negative, neutral, or positive – raises attention on all content published on the fan page. Therefore, the reach of negative eWOM increases with the volume of total eWOM. If both the volume of total and negative eWOM exceeds a certain threshold (cf. Phase 3), an online firestorm becomes likely. Figure 3 visualizes the time periods considered by the algorithm to predict the expected volume of overall and negative eWOM.



The second phase of the algorithm predicts the expected volume of overall and negative eWOM in the defined future time period t_0 . However, to yield reliable statistics and avoid rank deficiency within machine precision, the second phase of the algorithm is only executed if at least q percent of all C_i^{all} and C_i^{neg} are non-zero. Otherwise, the algorithm jumps directly to the third phase, and applies a default threshold to detect the emergence of an online firestorm (cf. Figure 2).

Phase 2: Predicting the Expected Volume of Overall and Negative eWOM

As the following steps are performed for overall and negative eWOM, we define $v \in \{all, neg\}$ and introduce a further variable $C_i^v \in \mathbb{N}_0 \forall i, v$ referring to either the overall or negative eWOM to simplify the notation. After determining the two time series C_i^v of overall and negative eWOM, we predict the expected volume \hat{C}_0^v for the future period t_0 . Therefore, we assume C_i^v to be Poisson distributed with mean $\mu^v \forall v$ and variance $\phi^v \mu^v \forall v$. This is in line with prior work predicting aspects of human activity (Malmgren et al. 2009), such as the purchase of consumer goods (Ehrenberg 1972; Fader et al. 2005; Schmittlein et al. 1987) and the spread of diseases (Farrington et al. 1996; Noufaily et al. 2013). The transferability of this concept to the context of online communication behavior should be discussed critically (Malmgren et al. 2009), as circadian cycles (that is, for example, some people's habit of only answering e-mails in the morning and the evening) have been reported to lead to heavy-tailed power-law distributions of inter-communication times (Malmgren et al. 2009). However, many studies have confirmed the applicability of a homogeneous Poisson process for modeling repeated communication behavior within these cycles, that is, for the above example, while being in one e-mail answering session (e.g., in the morning) (Malmgren et al. 2009). Because online firestorms emerge quickly and the window of opportunity for initiating countermeasures is very small (cf. previous section), only a relatively small number n of foregoing periods with a relatively short length Δt are considered when predicting the expected volume \hat{C}_0^v for the future period t_0 . In line with prior research (e.g., Probst 2011), we therefore assume that social media users are in one cycle while generating eWOM in the considered periods. Note that by taking only the direct foregoing periods into account, we implicitly consider seasonal variations in both volume and sentiment.

Based on the Poisson assumption, the algorithm identifies the most suitable Poisson model using the last n periods. The identified model is then used to forecast the overall number of eWOM as well as the number of negative eWOM for the future period t_0 , allowing the corresponding thresholds to be calculated. In line with Farrington et al. (1996) and Noufaily et al. (2013), we apply the following regression equation:

$$\log(E[C_i^v]) = \log(\mu^v) = \alpha^v + \beta^v * t_i \quad \forall i, \quad (1)$$

where α^v is a constant, β^v denotes the trend, and i is a control variable for all parameters of the last n periods. For an accurate prediction of the expected volume of eWOM \hat{C}_0^v in the future period t_0 , we must identify the most suitable Poisson regression model using maximum likelihood estimation. This selection procedure is iterative, and is performed two or four times in accordance with Noufaily et al. (2013) (cf.

Figure 2). The first, unweighted Poisson regression (i.e., the first iteration) is required to identify potential outliers, whose errors are down-weighted in the Maximum Likelihood estimation of the second, weighted Poisson regression (i.e., second iteration). This second iteration is used to examine whether the trend component β^v is significant (e.g., p-value < 0.05), and leads to a realistic prediction. In line with Noufaily et al. (2013), we assume every prediction to be unrealistic if it is smaller than $\min_{i=-n, \dots, -1} C_i^v$ or greater than $\max_{i=-n, \dots, -1} C_i^v$. In both cases (i.e., insignificant and/or unrealistic predictions), we perform a third, unweighted Poisson regression without a trend component (i.e., third iteration) to identify the outliers, which are again down-weighted in a weighted Poisson regression without considering the trend component (i.e., fourth iteration).

In the second and fourth iterations, the down-weighting procedure on the identified outliers is achieved with weights ω_i^v based on the Anscombe residuals (cf. Davison and Snell 1991), ensuring that values C_i^v with high residuals receive low weights and vice versa. To this end, the dispersion parameter ϕ^v is estimated as:

$$\hat{\phi}^v = \max \left\{ \frac{1}{n-p^v} \sum_{i=1}^n \omega_i^v \frac{(C_i^v - \hat{C}_i^v)^2}{\hat{C}_i^v}, 1 \right\}, \quad \hat{\phi}^v \geq 1, \quad (2)$$

where $p^v = \{1; 2\}$ denotes the degrees of freedom, which depend on whether a time trend is fitted (cf. Noufaily et al. 2013). The values of \hat{C}_i^v represent the fitted values, which are obtained by $\hat{C}_i^v = e^{\alpha^v + \beta^v * t_i}$. The weights ω_i^v are computed by:

$$\omega_i^v = \begin{cases} \gamma^v (s_i^v)^{-2}, & s_i^v > 1 \\ \gamma^v, & \text{otherwise} \end{cases}, \quad (3)$$

where γ^v is a constant such that $\sum_{i=-n}^{-1} \omega_i^v = n \forall v$, and s_i^v are the scaled Anscombe residuals, defined by:

$$s_i^v = \text{Anscombe Residuals} * \text{scalar} = \frac{3(C_i^{v^{2/3}} - \hat{C}_i^{v^{2/3}})}{2\hat{C}_i^{v^{1/6}}} * \frac{1}{\hat{\phi}^{v^{1/2}} (1 - h_{ii}^v)^{1/2}}, \quad (4)$$

where h_{ii}^v are the elements of the trace of the hat matrix.

The scaled Anscombe residuals, dispersion parameter, and weights are necessary to identify the most suitable Poisson model, from which the constant α^v and the trend component β^v can be inferred. The algorithm continues by calculating the expected volumes of overall and negative eWOM in the future period t_0 by applying the following formula:

$$\hat{C}_0^v = e^{\alpha^v + \beta^v * t_0}. \quad (5)$$

Phase 3: Calculating Thresholds and Checking whether to Raise an Alarm

If the actual volume of eWOM C_0^v in the future period t_0 exceeds the threshold T^v , the algorithm raises an alarm about the (statistical) emergence of an online firestorm.

To correct the skewness in the Poisson-distributed counts, we apply a 2/3 power transformation (cf. Farrington et al. 1996; Noufaily et al. 2013). Thus, we have an approximately symmetric distribution, and can derive accurate thresholds. The threshold T^v is an approximate $100(1 - \alpha)\%$ quantile for C_0^v , where z_α is the $100(1 - \alpha)$ percentile of the standard normal distribution with $\alpha \in (0; 1)$. However, if the check in the first phase has shown that less than q percent of all C_i^v are non-zero, the threshold is set to a default value $T_{\text{default}}^v \in \mathbb{R}_+$. This value is equal to the value from the period prior to t_0 . Hence, whenever it is not possible to determine (statistically) reliable thresholds, the previously calculated threshold remains valid. It could be argued that this increases the efficiency of the Online Firestorm Detector, as no update of previously calculated thresholds is required in periods of low activity. When the default threshold is required in the first iteration after the Online Firestorm Detector has been activated, we set it to the average eWOM level over the past n (micro-)periods, that is, $\frac{\sum_{i=-n}^{-1} C_i^v}{n}$. Taken together, the threshold T^v is defined as:

$$T^v = \begin{cases} \hat{C}_0^v \left[1 + 2/3 z_\alpha \frac{1}{\hat{C}_0^v} \left(\hat{\phi}^v \hat{C}_0^v + \text{var}(\hat{C}_0^v) \right) \right]^{3/2}, & \sum_{i=-n}^{-1} 1_{\{C_i^v > 0\}} \geq qn \\ T_{\text{default}}^v, & \text{otherwise} \end{cases}. \quad (6)$$

After both T^{all} and T^{neg} (for the overall and negative eWOM, respectively) have been calculated, the algorithm determines the actual overall volume C_0^{all} and negative volume C_0^{neg} of eWOM within the period t_0 to check whether the thresholds are exceeded. It is important to note that the alarm can be raised at any point in time, and not just at the end of every period Δt . As discussed before, our algorithm relies on both the volume of negative eWOM and the overall volume of eWOM to control for the level of attention on the fan page. This decreases the frequency of false alarms (type I errors), which might decrease the trustworthiness of the Online Firestorm Detector and lead to its warnings receiving less attention over time. However, the algorithm needs to be sensitive enough to detect online firestorms as soon as possible. Accounting for this trade-off, the algorithm raises an alarm if both the overall volume C_0^{all} and negative volume C_0^{neg} eWOM in the future period t_0 exceed the thresholds T^{all} and T^{neg} , respectively:

$$Alarm = \begin{cases} 1, & C_0^{all} > T^{all} \wedge C_0^{neg} > T^{neg} \\ 0, & \text{otherwise} \end{cases}, \quad (7)$$

where $Alarm \in \{0; 1\}$ denotes a dummy variable indicating whether an online firestorm alarm is raised ($Alarm = 1$) or not ($Alarm = 0$). At this stage, all responsible stakeholders (e.g., social media managers) could receive an alert via numerous possible channels (e.g., e-mail, SMS, phone-call).

Demonstration and Evaluation

In this section, we demonstrate the applicability of our proposed Online Firestorm Detector and evaluate its performance to “[...] measure how well the artifact supports a solution to the problem” (Peffer et al. 2007, p. 56). For this purpose, we implemented a prototype based on the software design paradigm of service-oriented architectures (cf. Papazoglou and Van Den Heuvel 2007). Hence, the prototype consists of several modules with clearly defined interfaces. Individual modules can be incrementally modified, maintained, or expanded if requirements or features change over time (e.g., to monitor different social media platforms or to conduct varying sentiment analyses). The three steps of the Online Firestorm Detector and the related implemented modules are illustrated in Figure 4.

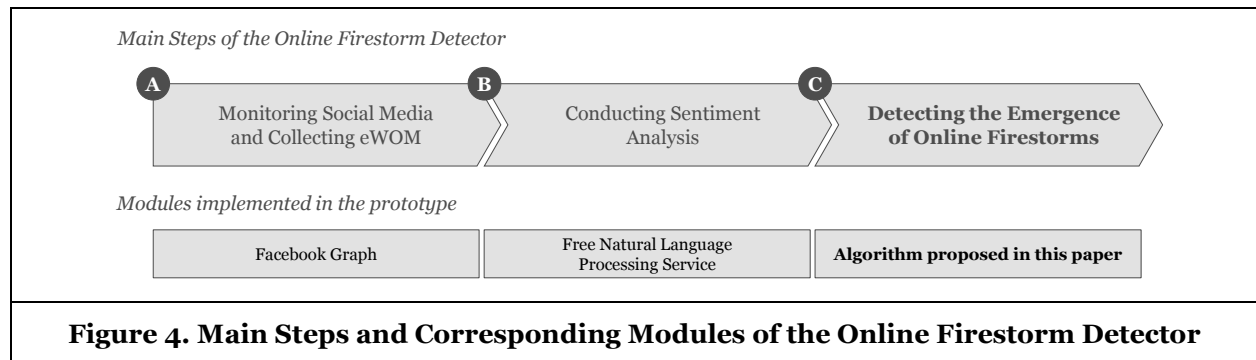


Figure 4. Main Steps and Corresponding Modules of the Online Firestorm Detector

To empirically evaluate the proposed approach, we require a past online firestorm that meets three criteria: First, the outburst of negative eWOM must be in line with the definition of online firestorms provided by Pfeffer et al. (2014). Second, the online firestorm should be large enough or have sufficient press coverage (i.e., reporting in channels outside of social media) to substantiate its impact. Third, the social media presence of the company had to be (generally) highly penetrated by eWOM, so that manual surveillance was inapplicable, and automated, real-time detection was therefore reasonable. Considering all three criteria, the authors manually investigated news portals and mutually agreed that the online firestorm suffered by Coca-Cola on its Facebook page (/cocacola) was suited to our purposes. Additionally, we examined the reason for the online firestorm (controversial commercial featuring “America the Beautiful” during Super Bowl XLVIII) and the point at which it started (i.e., time when the first related negative eWOM was created: February 3rd 2014, 00:50 GMT). For demonstration and evaluation, we consider the three months before (“pre-firestorm-era”) and one month after the beginning (“post-firestorm-era”). The pre-firestorm-era allows the frequency of false alarms (type I errors) to be evaluated, and the post-firestorm-era allows the duration until the online firestorm was detected to be measured.

Step (A): Monitoring Social Media and Collecting eWOM

As outlined above, the Online Firestorm Detector must first monitor and collect company-related eWOM generated across social media in real time. Therefore, the prototype implemented for the demonstration and evaluation of the Online Firestorm Detector contains a module that allows eWOM to be retrieved from Facebook, as this is a particularly important source of eWOM (cf. e.g., Berger et al. 2014; Heidemann et al. 2012). In general, the module was implemented in Java using the external library “RestFB” for connecting to the Facebook API, authorizing (via “OAuth”), and processing the returned data (“JSON” format). As a result, a unique ID, timestamp, and the content of each eWOM (i.e., the message) can be extracted (i.e., wallposts and related comments from a company’s fan page). Thus, the module allows for the extraction of all past data and the continuous collection of wallposts and comments (i.e., eWOM) from a company’s Facebook fan page (in this case, /cocacola). Table 2 presents the number of overall eWOM for both the (three-month) pre-firestorm-era and the (one-month) post-firestorm-era. The daily mean and standard deviation (SD) for both eras are also shown. It can be observed that the mean and standard deviation rise considerably after the beginning of the online firestorm (cf. Table 2).

Table 2. Summary Statistics for Step (A) of the Online Firestorm Detector			
Number of Overall eWOM		Mean (SD) of Overall eWOM per Day	
Pre-firestorm-era	Post-firestorm-era	Pre-firestorm-era	Post-firestorm-era
51,909	27,669	559.5 (664.4)	936.8 (1080.8)

Step (B): Conducting Sentiment Analysis

Step (B) complements the existing unique ID, timestamp, and content of each piece of eWOM by a flag indicating whether the eWOM is of negative sentiment (cf. previous section). For the prototype implementation, we created a module in Java utilizing the API of the Free Natural Language Processing Service (loudelement.com), which allows eWOM to be classified as “negative,” “neutral,” or “positive.” We chose this particular provider based on four requirements. First, whole text and not only single words should be classified. Second, multiple languages should be supported. Third, the service had to be free of cost for the prototype to meet budget constraints. Fourth, the reliability of the sentiment analysis had to be sufficient. To assess the quality of sentiment analysis tools, evaluation datasets can be used (cf. e.g., Saif et al. 2013). Therefore, we applied the Free Natural Language Processing Service to a manually classified dataset that has been used in previous research on eWOM in social media (cf. Scholz et al. 2013). The quality of the resulting classification was assessed as being sufficient. As we are only interested in flagging negative eWOM, we did not distinguish between neutral and positive sentiment. While the sentiment analysis provider executes the classification, the module is able to receive eWOM collected in step (A) and to serve step (C) with the classified results (negative or not negative) via defined interfaces. In the future, further (proprietary) sentiment analysis tools (e.g., SPSS Clementine) could be integrated to provide multiple sentiment analysis results, which would help to minimize the likelihood of incorrect classifications. We applied the module to the eWOM collected from the Coca-Cola Facebook fan page to analyze its sentiment. As we can see from Table 3, about 9% of the overall eWOM from the pre-firestorm-era is classified as negative. In the post-firestorm-era, this number increases to over 16%. Furthermore, we can observe that both the mean and standard deviation of the daily negative eWOM increase.

Table 3. Summary Statistics for Step (B) of Online Firestorm Detector			
Number of Negative eWOM (% of overall eWOM)		Mean (SD) of Negative eWOM per Day	
Pre-firestorm-era	Post-firestorm-era	Pre-firestorm-era	Post-firestorm-era
4,551 (8.77%)	4,520 (16.33%)	49 (23.2)	152 (382.1)

Step (C): Detecting the Emergence of Online Firestorms

The actual calculation in step (C) was implemented within MATLAB. Before applying the Online Firestorm Detector, we had to initialize its three main parameters:

Length of (micro-)period Δt : A necessary condition for the early detection of an online firestorm is the continuous surveillance of ongoing activity in terms of generated eWOM. However, there is a trade-off between shorter and longer (micro-)periods t_i . On the one hand, the shorter the length of Δt , the better the consideration of short-term variations in both overall and negative eWOM. On the other hand, the corresponding overall and negative eWOM counts C_i^v become (on average) very small, leading to lower default thresholds and a higher volatility between the (micro-)periods' overall and negative eWOM. In interviews with experts from business practice, the average time from the detection of an online firestorm until effective countermeasures must be implemented has been identified as 7 min. Therefore, we selected a lower default value of $\Delta t = 5$ min. However, to evaluate the Online Firestorm Detector's robustness, we also applied $\Delta t = 15$ min and $\Delta t = 60$ min.

Number of (micro-)periods n : There is also a trade-off when choosing the number of (micro-)periods n . On the one hand, a prediction based on fewer periods may consider seasonality, but more periods may improve the statistical explanatory power of the prediction. Considering this, we selected a default value of $n = 30$, and also applied $n = 15$ and $n = 45$ to verify the Online Firestorm Detector's robustness.

Probability of error α : The third parameter to be determined is the probability of error α . Again, although a lower probability of error reduces false alarms, there is a trade-off because this will lower the probability of detecting the emergence of an online firestorm. As suggested in the literature on epidemiological surveillance, we applied a default value of $\alpha = 0.05$ (Freeman et al. 2013; Noufaily et al. 2013). We varied this parameter to $\alpha = 0.01$ and $\alpha = 0.10$, which are common levels of significance.

Besides these three main parameters, we must determine a value for q , that is, the percentage of zeros that determines whether a new threshold is calculated. We set $q = 25\%$ to increase the likelihood of nonsingular and well-conditioned matrices, and to increase the Online Firestorm Detector's efficiency. As variations in q had little effect on the Online Firestorm Detector's results, we do not include them here. However, the authors are happy to share these results upon request. Based on the collected and classified eWOM, the Online Firestorm Detector executed the three phases of the detection algorithm (cf. Figure 2). In the first phase, that is, the preprocessing of the data, the algorithm counts the number of overall and negative eWOM generated during the last n (micro-)periods. As outlined in the previous section, the algorithm applies the default thresholds $T_{\text{default}}^{\text{all}}$ and $T_{\text{default}}^{\text{neg}}$ (i.e., the previously calculated thresholds) if the activity in the previous periods has been too low (here: in less than 25% of the n last (micro-)periods any eWOM has been generated) to determine (statistically) reliable new thresholds for the future period t_0 . This helps to increase the efficiency of the algorithm, as it removes the need to fit the Poisson models in periods of low activity, thus reducing the computation time. In the second phase, the algorithm predicts the expected volume of overall and negative eWOM by applying Formulas (1)–(5). In the third phase, new thresholds are calculated based on Formula (6). Furthermore, the Online Firestorm Detector continuously determines the actual volume of overall C_0^{all} and negative C_0^{neg} eWOM within the future period t_0 to check whether the thresholds are likely to be exceeded. Based on Formula (7), the algorithm raises an alarm as soon as both the volume of overall C_0^{all} and negative C_0^{neg} eWOM reach the respective thresholds T^{all} and T^{neg} .

As shown in Table 4, regardless of the chosen parameters (Δt , n , and α), our proposed Online Firestorm Detector reliably detected the outbreak of Coca-Cola's online firestorm 46–79 s after the first firestorm-related eWOM. In the default setting, the online firestorm was detected within 60 s. The number of false alarms (type I errors) was low (from 1–17 during the three months of the pre-firestorm-era, cf. Table 4). In the default setting, only one false alarm was raised before the actual outburst after three months of observation. Regarding the Online Firestorm Detector's sensibility to variations in the parameter settings, Table 4 shows that variations in the length of (micro-)periods Δt and in the number of (micro-)periods n had a very small influence on the time until the online firestorm was detected. At the same time, the number of false alarms increased with the length and number of (micro-)periods, as the thresholds were less frequently adapted and variations of eWOM over time were less well accounted for. With respect to the probability of error α , we can observe the intuitively expected behavior of a reduction in detection time and an increasing number of false alarms as the probability of error increased. Taken together, the variations in

the parameters summarized in Table 4 show that the Online Firestorm Detector is robust to parameter changes and its results are reliable.

		<i>n</i> = 15			<i>n</i> = 30			<i>n</i> = 45		
		$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$	$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$	$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$
$\Delta t = 5$ min	time elapsed ¹	60	54	46	71	60	54	72	60	54
	type I errors ²	1	2	5	1	1	3	1	1	2
$\Delta t = 15$ min	time elapsed ¹	60	54	46	72	60	54	73	60	54
	type I errors ²	2	6	17	2	3	11	2	2	7
$\Delta t = 60$ min	time elapsed ¹	73	60	54	73	60	60	79	71	60
	type I errors ²	5	18	63	5	11	36	5	10	22

¹ Duration between first firestorm-related eWOM and alarm in seconds
² Number of false alarms during the pre-firestorm-era
 Online Firestorm Detector's default parameters and related results highlighted in grey

Competitive Benchmarking

As mentioned before, we know of no alternative approaches for the detection of online firestorms from prior research (so-called “competing artifacts”, Hevner et al. 2004). Thus, a comparison against existing scientific approaches is not possible. Instead, we evaluate the Online Firestorm Detector against two potential rules of thumb (i.e., thresholds based on the mean and maximum of eWOM within the last *n* periods), as well as against a straightforward linear regression estimation. To underpin the need for both the sentiment analysis in step (B) and the decision to consider negative as well as total eWOM within step (C), we provide two further internal benchmarks by varying our Online Firestorm Detector.

First, it may appear appropriate to rely solely on descriptive statistics of past negative eWOM, for example, using the mean or maximum value of past eWOM as a threshold and triggering an alarm as soon as the actual number of negative eWOM exceeds the respective mean or maximum from the last period. To allow for a fair comparison, we decided to keep the underlying parameters (i.e., Δt and *n*) unchanged. Looking at Table 5, regardless of variations in Δt and *n*, both the mean and maximum approach detect the outbreak of the online firestorm within 60 s. However, the several hundred false alarms (cf. Table 5) are not acceptable. Thus, considering the crucial trade-off between detection time and false alarms (cf. e.g., Wagner et al. 2001), the Online Firestorm Detector performs considerably better than common rules of thumb that could alternatively be used.

		<i>n</i> = 15		<i>n</i> = 30		<i>n</i> = 45	
		Mean	Max	Mean	Max	Mean	Max
$\Delta t = 5$ min	time elapsed ¹	30	31	30	31	30	31
	type I errors ²	2,464	395	2,457	223	2,456	160
$\Delta t = 15$ min	time elapsed ¹	30	32	30	46	30	46
	type I errors ²	1,421	313	1,403	182	1,388	137
$\Delta t = 60$ min	time elapsed ¹	31	60	30	60	30	60
	type I errors ²	295	154	298	89	285	69

¹ Duration between first firestorm-related eWOM and alarm in seconds
² Number of false alarms during the pre-firestorm-era

Second, to challenge the Online Firestorm Detector, we used linear regression estimation instead of the adapted (iterative) Poisson model. To enable a fair comparison, we again decided to keep the common parameters (i.e., Δt , *n*, and α) unchanged. As depicted in Table 6, the online firestorm was detected

within 79 s, although up to 673 false alarms were raised. Considering the crucial trade-off between detection time and false alarms (cf. e.g., Wagner et al. 2001), the Online Firestorm Detector therefore performs considerably better than a linear regression procedure that could alternatively be used.

		$n = 15$			$n = 30$			$n = 45$		
		$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$	$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$	$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$
$\Delta t = 5$ min	time elapsed ¹	30	30	30	30	30	30	30	30	30
	type I errors ²	77	279	544	108	369	624	135	393	673
$\Delta t = 15$ min	time elapsed ¹	31	31	30	32	31	31	32	31	31
	type I errors ²	36	165	319	54	190	344	66	207	356
$\Delta t = 60$ min	time elapsed ¹	79	60	46	71	46	40	71	46	40
	type I errors ²	25	77	142	23	74	124	19	61	111

¹ Duration between first firestorm-related eWOM and alarm in seconds
² Number of false alarms during the pre-firestorm-era

Third, it may seem feasible to skip the sentiment analysis within step (B) and rely on (unclassified) total eWOM within the novel algorithm utilized in step (C). As shown in Table 7, with the same common parameter values (i.e., Δt , n , and α), both the detection time and the number of false alarms increased considerably. These results emphasize the need to consider the sentiment of eWOM, and thus the sentiment analysis within step (B).

		$n = 15$			$n = 30$			$n = 45$		
		$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$	$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$	$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$
$\Delta t = 5$ min	time elapsed ¹	73	60	54	60	71	83	60	71	83
	type I errors ²	29	140	409	24	73	209	20	47	139
$\Delta t = 15$ min	time elapsed ¹	90	79	71	117	83	79	127	90	80
	type I errors ²	48	188	415	32	110	259	27	77	192
$\Delta t = 60$ min	time elapsed ¹	183	166	157	171	163	155	200	169	163
	type I errors ²	82	181	237	46	118	168	32	89	152

¹ Duration between first firestorm-related eWOM and alarm in seconds
² Number of false alarms during the pre-firestorm-era

We also benchmarked the proposed approach against an adapted approach that considers only negative eWOM. Again, the common parameters (i.e., Δt , n , and α) were kept the same. As shown in Table 8, the detection times were similar to those in the proposed approach. However, the number of false alarms generally increased. Even though these results might look promising at first sight, those from the proposed approach were significantly better in terms of false alarms (p-value: 0.014).

Table 8. Summary Statistics Considering Solely Negative eWOM

		n = 15			n = 30			n = 45		
		$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$	$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$	$\alpha = 0.01$	$\alpha = 0.05$	$\alpha = 0.10$
$\Delta t = 5 \text{ min}$	time elapsed ¹	60	46	45	70	54	46	71	54	46
	type I errors ²	1	2	6	1	2	4	1	1	4
$\Delta t = 15 \text{ min}$	time elapsed ¹	60	46	45	70	54	46	73	60	46
	type I errors ²	3	3	25	3	3	13	3	3	7
$\Delta t = 60 \text{ min}$	time elapsed ¹	79	70	54	80	70	60	83	71	60
	type I errors ²	7	25	59	6	14	35	5	11	27

¹ Duration between first firestorm-related eWOM and alarm in seconds
² Number of false alarms during the pre-firestorm-era

Discussion on Contribution, Limitations, and Further Research

Contribution to Theory and Practice

As shown in the previous section, the proposed Online Firestorm Detector (“design artifact”) can reliably detect the emergence of online firestorms shortly after the first piece of related negative eWOM has been generated. Furthermore, the number of false alarms (type I errors) is low. The demonstration and evaluation hence shows that the proposed and prototypically implemented Online Firestorm Detector “[...] works and does what it is meant to do” (“validity”, Gregor and Hevner 2013, p. 351). By using real-world data directly collected from a Facebook fan page, we also demonstrated the usability of our artifact in business practice (“utility”, Hevner et al. 2004). By means of competitive benchmarking, we were moreover able to show that the proposed Online Firestorm Detector outperforms competing artifacts (“quality”, Gregor and Hevner 2013). The design, demonstration, and evaluation of our Online Firestorm Detector contribute to both theory and practice. From a theoretical perspective, our contribution to the literature on IS and marketing is threefold:

First, we enrich existing IS and marketing literature on the analysis of eWOM in social media to avert its potential dark side: Existing IS and marketing literature on the analysis of eWOM in social media has mainly focused on the positive aspects (cf. Berger et al. 2014). For instance, multiple studies investigated the diffusion of positive eWOM through social networks to increase marketing efficiency by targeting a set of influential users (cf. Probst et al. 2013). However, research focusing on the potential negative consequences of companies’ social media engagement (“dark side of social media”), and particularly negative eWOM, is rare (cf. Probst et al. 2013). Although a few studies deal with the dynamics (Pfeffer et al. 2014; Stich et al. 2014) or restriction of online firestorms (Mochalova and Nanopoulos 2014; Munzel et al. 2012), the question of when to trigger an alarm if negative eWOM spreads over an entire network has not yet been addressed. Hence, the design of our Online Firestorm Detector constitutes an essential element in averting the potential negative consequences of companies’ social media engagements by promoting research on early detection approaches for online firestorms.

Second, we have contributed to a valid theoretical basis for research on eWOM diffusion in social media: We argued that epidemics and online firestorms share commonalities, and that approaches for combating epidemics may deductively be adapted to combat online firestorms. The successful empirical demonstration and evaluation of our artifact shows that research findings regarding the early detection of outbreaks of infectious disease may indeed be exapted to the context of social media. Hence, this study supports earlier work which argued that the spread of negative eWOM can be compared to the spreading behavior of pathogenic organisms (cf. e.g., Budak et al. 2011; Kempe et al. 2003). Furthermore, our results encourage the use of regression model-based techniques (cf. Chandola et al. 2009) and the adoption of existing approaches from epidemiological surveillance systems relying on so-called anomaly scores (cf. Anscombe and Guttman 1960) to identify contextual anomalies in social media. Thus, we have demonstrated that

research from the field of epidemiology can serve as a valid theoretical basis for design science research in the context of eWOM diffusion in social media.

Third, we have extended the understanding of online firestorms by showing that not only negative eWOM should be considered in their detection: In the literature, online firestorms are defined as an “[...] sudden discharge of [...] negative [e]WOM” (Pfeffer et al. 2014, p. 118). However, our empirical demonstration and evaluation has shown that the design of the Online Firestorm Detector in considering both negative and total eWOM leads to significantly less false alarms than alternative implementations relying solely on negative eWOM. This is particularly important, as the competitive benchmarking with common rules of thumb and linear regression has shown that it is rather straightforward to achieve short detection times but difficult to avoid a high error rate. Moreover, we must emphasize the importance of reliable sentiment classification techniques, as alternative implementations building solely on the total volume of eWOM performed considerably worse in terms of both detection time and error rate.

Alongside these theoretical contributions, we have also contributed to business practice.

First, we have shown that common lightweight solutions are unable to reliably detect online firestorms in social media: By means of competitive benchmarking, we have demonstrated that mean or maximum rules, or even apparently more sophisticated procedures such as linear regression estimation, deliver insufficient reliability, that is, too many false alarms. Consequently, unnecessary resources would be needed to verify whether an alarm is actually correct. Moreover, confidence, trust, and attention would quickly dissolve, as the behavioral intention to use technology depends, among other things, on “performance expectancy”, that is, “[...] the degree to which a person believes that using a particular system would enhance his or her job performance” (Venkatesh et al. 2003, p. 447). Our suggested Online Firestorm Detector represents a well-founded approach for the automated, real-time detection of emerging online firestorms in social media that allows for timely detection and a small number of false alarms.

Second, we have provided a ready-to-use artifact that enables companies to mitigate risks from social media engagements: To prevent a domino effect and avoid a steep drop in customer satisfaction ratings or even share prices (cf. Dell Hell), timely and reliable detection is crucial to the initiation of countermeasures. Thus, this study contributes to business practice by both emphasizing the importance of an automated real-time detection system and providing a ready-to-use artifact. The suggested Online Firestorm Detector allows high-profile, hard-to-predict, and rare events such as online firestorms to be detected, thus enabling risk mitigation within organizations. This is particularly valuable for large companies that have widely adopted eWOM marketing in social media, as well as small companies that lack dedicated resources for extensive social media monitoring.

Limitations and Further Research

As well as the promising results described above, there are limitations that provide room for improvement and starting points for further research.

First, we evaluated our artifact based on a single firestorm suffered by Coca-Cola on Facebook. Thus, the detection of online firestorms in other social media outlets has not yet been assessed. However, we intend to implement further modules that allow access to company-related eWOM generated in other social media such as micro-blogging and blogs. Furthermore, we focused on one online firestorm that could be manually and ex post detected within Facebook. Therefore, even though we successfully proved the reliability of our artifact by varying certain parameters, its generalizability remains to be confirmed in future research (e.g., by analyzing different firestorms or different social media).

Second, the parameterization of the number and length of the (micro-)periods needs to be critically examined. Based on the underlying trade-offs discussed in this paper, we chose default parameters in line with statistical practice (regarding the number of (micro-)periods) and based on the findings of expert interviews (regarding the length of (micro-)periods). Even though we generalized our results by comparing the output from different parameterizations, and these variations in the parameters did not lead to unexpected deviations, future research could improve the parameter settings by training the algorithm on past online firestorms related to comparable companies (e.g., approximated in terms of mean and volatility of eWOM, homophily of customers).

Third, sentiment analysis of eWOM in social media is challenging for both practitioners and IS researchers because of personal, cultural, and contextual factors such as irony (cf. Feldman 2013; Pang and Lee 2008). Although we assessed the quality of our sentiment analysis using an evaluation dataset from previous research (cf. Scholz et al. 2013), the automated classification of eWOM may potentially differ among humans. To enable an easy exchange of the sentiment analysis service when newer or better algorithms emerge, we deliberately designed our Online Firestorm Detector based on the software design paradigm of service-oriented architectures (cf. Papazoglou and Van Den Heuvel 2007). We suggest that varying sentiment analysis services should be used in future research to compare their results.

Fourth, we did not explicitly address potential interaction effects between negative, neutral, and positive eWOM. For example, there might be situations, where a positive reaction of the fan base on negative eWOM (leading to a strong increase in positive eWOM) avoids the outburst of an online firestorm. However, such interaction effects have not been sufficiently empirically studied yet and are beyond the scope of this paper. As we discussed before, the Online Firestorm Detector's reliability is of utmost importance; so is a low number of false alarms. One could argue that our current design is very sensitive, because we raise an alarm even though there might be enough positive eWOM to stop the diffusion of negative eWOM. Due to the low number of false alarms, however, we believe that neglecting potential interaction effects is reasonable in a first step. Future work is encouraged to study the interplay of negative, neutral, and positive eWOM.

Finally, we did not address the question of how to deal with online firestorms after their emergence has been detected. Though this is beyond the scope of the current paper, we encourage research that builds on existing studies (Mochalova and Nanopoulos 2014; Munzel et al. 2012) to develop strategies for the actual mitigation of online firestorms. Moreover, the organizational acceptance of artifacts such as our Online Firestorm Detector in dependence of their reliability (here: false alarms) should be studied in future IS research.

Conclusion

In line with the publication schema for Design Science Research proposed by Gregor and Hevner (2013), we can summarize as follows: First, we specified the “*purpose and scope*” of our artifact, stated its “*relevance*” for business practice, and identified the absence of approaches for an automated, real-time detection of online firestorms as research gap. Second, we emphasized that we address an “*important and relevant business problem*” (Hevner et al. 2004, p. 83) by presenting prior empirical research, which provided evidence that negative eWOM leads to negative economic effects. Moreover, we discussed “*descriptive and prescriptive knowledge*” by discussing related work on the diffusion of eWOM as well as the context of anomaly detection and epidemiological surveillance, which informed the design of our artifact. Third, we developed our “*design artifact*”, the Online Firestorm Detector, which has led to two design science research contributions (cf. Gregor and Hevner 2013, p. 342): As the first contribution, we designed an algorithm for the detection of emerging online firestorms (i.e., “*level 2 research contribution*”), which was inspired by prior work from the field of epidemiological surveillance systems (i.e., “*knowledge contribution by exaptation*”). To support a “*rigorous*” definition and presentation, we formally denoted the algorithm (Hevner et al. 2004). As the second contribution, we implemented a prototype of the Online Firestorm Detector (i.e., “*level 1 research contribution*”), comprising three main steps: (A) Monitoring social media and collecting eWOM, (B) conducting sentiment analysis, and (C) detecting the emergence of online firestorms by applying the abovementioned algorithm. Subsequently, we provided a “*rigorous design evaluation*” to demonstrate the validity and utility of the artifact by applying the Online Firestorm Detector to the situation suffered by Coca-Cola on Facebook, thus evaluating it “*in depth in business*” (Hevner et al. 2004). We thereby showed that the Online Firestorm Detector's algorithm reliably detects emerging online firestorms shortly after the emergence of the first piece of related negative eWOM, and that the number of false alarms is low. We furthermore provided empirical evidence that our design artifact is superior to potential competing artifacts, such as common rules of thumb (mean and maximum) or linear regression estimation. Finally, we critically “*discussed and reflected*” on our artifact by pointing out its theoretical and practical contributions, limitations, and areas for further research. We are confident that a further developed prototype can soon be applied in business practice, and that our Online Firestorm Detector already serves as a sound starting point for future research.

Acknowledgements

This research was (in part) carried out in the context of the Project Group Business and Information Systems Engineering of the Fraunhofer Institute for Applied Information Technology FIT. Furthermore, the authors would like to thank Nicholas Berente and Henner Gimpel for their helpful comments.

References

- Agarwal, D. 2005. "An Empirical Bayes Approach to Detect Anomalies in Dynamic Multidimensional Arrays," in *Proceedings of the 5th IEEE International Conference on Data Mining*, Houston, USA, pp. 26–33.
- Albuquerque, P., Pavlidis, P., Chatow, U., Chen, K.-Y., and Jamal, Z. 2012. "Evaluating Promotional Activities in an Online Two-Sided Market of User-Generated Content," *Marketing Science* (31:3), pp. 406–432.
- Anderson, E. W. 1998. "Customer Satisfaction and Word of Mouth," *Journal of Service Research* (1:1), pp. 5–17.
- Anscombe, F. J., and Guttman, I. 1960. "Rejection of Outliers," *Technometrics* (2:2), pp. 123–147.
- Bampo, M., Ewing, M. T., Mather, D. R., Stewart, D., and Wallace, M. 2008. "The Effects of the Social Structure of Digital Networks on Viral Marketing Performance," *Information Systems Research* (19:3), pp. 273–290.
- Bansal, S., Grenfell, B. T., and Meyers, L. A. 2007. "When individual behaviour matters: homogeneous and network models in epidemiology," *Journal of the Royal Society Interface* (4:16), pp. 879–891.
- Banz, M., Grosswiele, L., Huber, J., and Probst, F. 2015. "More Fans at any Cost? Analyzing the Economic Effects of the Ratio of Fans to Non-Fans in a Customer Portfolio Considering Electronic Word-of-Mouth," *Journal of Decision Systems* (accepted with minor revisions).
- Benevenuto, F., Rodrigues, T., Cha, M., and Almeida, V. 2009. "Characterizing user behavior in online social networks," in *Proceedings of the 9th ACM SIGCOMM Conference*, Chicago, USA, pp. 49–62.
- Berger, J., and Milkman, K. L. 2012. "What Makes Online Content Viral?" *Journal of Marketing Research* (49:2), pp. 192–205.
- Berger, K., Klier, J., Klier, M., and Probst, F. 2014. "A Review of Information Systems Research on Online Social Networks," *Communications of the Association for Information Systems* (35:1), article 8.
- Brown, R. G. 1971. "Detection of turning points in a time series," *Decision Sciences* (2:4), pp. 383–403.
- Budak, C., Agrawal, D., and Abbadi, A. E. 2011. "Limiting the Spread of Misinformation in Social Networks," in *Proceedings of the 20th International Conference on World Wide Web*, Hyderabad, India, pp. 665–674.
- Chandola, V., Banerjee, A., and Kumar, V. 2009. "Anomaly Detection: A Survey," *ACM Computing Surveys* (41:3), pp. 1–72.
- Chen, P.-Y., Wu, S.-y., and Yoon, J. 2004. "The Impact of Online Recommendations and Consumer Feedback on Sales," in *Proceedings of the 25th International Conference on Information Systems*, Washington DC, USA, paper 58.
- Chen, W., Collins, A., Cummings, R., Ke, T., Liu, Z., Rincon, D., Sun, X., Wang, Y., Wei, W., and Yuan, Y. 2011. "Influence Maximization in Social Networks When Negative Opinions May Emerge and Propagate," in *Proceedings of the 11th SIAM International Conference on Data Mining*, Philadelphia, USA, pp. 379–390.
- Chen, W., Wang, Y., and Yang, S. 2009. "Efficient influence maximization in social networks," in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Paris, France, pp. 199–208.
- Chen, Y., and Xie, J. 2008. "Online Consumer Review: Word-of-Mouth as a New Element of Marketing Communication Mix," *Management Science* (54:3), pp. 477–491.
- Chevalier, J. A., and Mayzlin, D. 2006. "The Effect of Word of Mouth on Sales Online Book Reviews," *Journal of Marketing Research* (43:3), pp. 345–354.
- Chintagunta, P. K., Gopinath, S., and Venkataraman, S. 2010. "The Effects of Online User Reviews on Movie Box Office Performance Accounting for Sequential," *Marketing Science* (29:5), pp. 944–957.
- Davison, A. C., and Snell, E. J. 1991. "Residuals and Diagnostics," in *Statistical Theory and Modelling*, D. V. Hinkley, N. Reid and E. J. Snell (eds.), London: Chapman & Hall.

- Dellarocas, C. 2003. "The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms," *Management Science* (49:10), pp. 1407–1424.
- Dellarocas, C., Zhang, X., and Awad, N. F. 2007. "Exploring the value of online product reviews in forecasting sales The case of motion pictures," *Journal of Interactive Marketing* (21:4), pp. 23–45.
- Dhar, V., and Chang, E. A. 2009. "Does Chatter Matter? The Impact of User-Generated Content on Music Sales," *Journal of Interactive Marketing* (23:4), pp. 300–307.
- Domingos, P., and Richardson, M. 2001. "Mining the Network Value of Customers," in *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, USA, pp. 57–66.
- Duan, W., Gu, B., and Whinston, A. B. 2008. "Do online reviews matter?—An empirical investigation of panel data," *Decision support systems* (45:4), pp. 1007–1016.
- Ehrenberg, A. S. 1972. *Repeat-buying; theory and applications*, New York: American Elsevier.
- Faase, R., Helms, R., and Spruit, M. 2011. "Web 2.0 in the CRM domain: defining social CRM," *International Journal of Electronic Customer Relationship Management* (5:1), pp. 1–22.
- Fader, P. S., Hardie, Bruce G. S., and Lee, K. L. 2005. "'Counting Your Customers' the Easy Way: An Alternative to the Pareto/NBD Model," *Marketing Science* (24:2), pp. 275–284.
- Farrington, C. P., Andrews, N. J., Beale, A. D., and Catchpole, M. A. 1996. "A Statistical Algorithm for the Early Detection of Outbreaks of Infectious Disease," *Journal of the Royal Statistical Society* (159:3), pp. 547–563.
- Feldman, R. 2013. "Techniques and applications for sentiment analysis," *Communications of the ACM* (56:4), p. 82.
- Forman, C., Ghose, A., and Wiesenfeld, B. 2008. "Examining the Relationship Between Reviews and Sales The Role of Reviewer Identity Disclosure in Electronic Markets," *Information Systems Research* (19:3), pp. 291–313.
- Fox, A. J. 1972. "Outliers in Time Series," *Journal of the Royal Statistical Society* (Series B, 34), pp. 350–363.
- Freeman, R., Charlett, A., Hopkins, S., Am O'Connell, Andrews, N., Freed, J., Holmes, A., and Catchpole, M. 2013. "Evaluation of a national microbiological surveillance system to inform automated outbreak detection," *Journal of Infection* (67:5), pp. 378–384.
- Furfie, B. 2008. *Dell Hell' PR Disaster Could Take Ten Years to Fix*. <http://www.pcr-online.biz/news/read/dell-hell-pr-disaster-could-take-ten-years-to-fix/>. Accessed 2 October 2013.
- Goh, K.-Y., Heng, C.-S., and Lin, Z. 2013. "Social Media Brand Community and Consumer Behavior: Quantifying the Relative Impact of User- and Marketer-Generated Content," *Information Systems Research* (24:1), pp. 88–107.
- Gregor, S., and Hevner, A. R. 2013. "Positioning and Presenting Design Science Research for Maximum Impact," *MIS Quarterly* (37:2), pp. 337–355.
- Guillou, A., Kratz, M., and Le Strat, Y. 2014. "An extreme value theory approach for the early detection of time clusters. A simulation-based assessment and an illustration to the surveillance of Salmonella," *Statistics in medicine* (33:28), pp. 5015–5027.
- Harris, L., and Dennis, C. 2011. "Engaging Customers on Facebook: Challenges for e-Retailers," *Journal of Consumer Behaviour* (10:6), pp. 338–346.
- He, X., Song, G., Chen, W., and Jiang, Q. 2012. "Influence Blocking Maximization in Social Networks under the Competitive Linear Threshold Model," in *Proceedings of the 12th SIAM International Conference on Data Mining*, Anaheim, USA, pp. 463–474.
- Heidemann, J., Klier, M., and Probst, F. 2012. "Online social networks: A survey of a global phenomenon," *Computer Networks* (56:18), pp. 3866–3878.
- Hennig-Thurau, T., Edward, M. C., Friege, C., Gensler, S., Lobschat, L., Rangaswamy, A., and Skiera, B. 2010. "The Impact of New Media on Customer Relationships," *Journal of Service Research* (13:3), pp. 311–330.
- Hennig-Thurau, T., Gwinner, K. P., Walsh, G., and Gremler, D. D. 2004. "Electronic word-of-mouth via consumer-opinion platforms," *Journal of Interactive Marketing* (18:1), pp. 38–52.
- Hevner, A. R., March, S. T., Park, J., and Ram, S. 2004. "Design Science in Information Systems Research," *MIS Quarterly* (28:1), pp. 75–105.
- Horn, P. S., Feng, L., Li, Y., and Pesce, A. J. 2001. "Effect of outliers and nonhealthy individuals on reference interval estimation," *Clinical Chemistry* (47:12), pp. 2137–2145.

- Hulth, A., Andrews, N., Ethelberg, S., Dreesman, J., Faensen, D., van Pelt, W., and Schnitzler, J. 2010. "Practical usage of computer-supported outbreak detection in five European countries," *Eurosurveillance* (15:36), pp. 1–6.
- Jarvis, J. 2005. *My Dell Hell*. <http://www.theguardian.com/technology/2005/aug/29/mondaymediasection.blogging>. Accessed 30 November 2013.
- Kaplan, A. M., and Haenlein, M. 2010. "Users of the world, unite! The challenges and opportunities of Social Media," *Business Horizons* (53:1), pp. 59–68.
- Kay, S. M. 1998. *Fundamentals of statistical signal processing*, Englewood Cliffs, USA: Prentice-Hall.
- Kempe, D., Kleinberg, J., and Tardos, É. 2003. "Maximizing the spread of influence through a social network," in *Proceedings of 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, New York, USA, pp. 137–146.
- Kennedy, H. 2012. "Perspectives on Sentiment Analysis," *Journal of Broadcasting & Electronic Media* (56:4), pp. 435–450.
- Kietzmann, J. H., Hermkens, K., McCarthy, I. P., and Silvestre, B. S. 2011. "Social media? Get serious! Understanding the functional building blocks of social media," *Business Horizons* (54:3), pp. 241–251.
- Laurikkala, J., Juhola, M., and Kentala, E. 2000. "Informal identification of outliers in medical data," in *Proceedings of the 5th International Workshop on Intelligent Data Analysis in Medicine and Pharmacology*, Berlin, Germany, pp. 24–28.
- Leskovec, J., Adamic, L. A., and Huberman, B. A. 2007. "The dynamics of viral marketing," *ACM Transactions on the Web (TWEB)* (1:1), article 5.
- Libai, B., Bolton, R., Bügel, M. S., Ruyter, K. de, Götz, O., Risselada, H., and Stephan, A. T. 2010. "Customer-to-Customer Interactions: Broadening the Scope of Word of Mouth Research," *Journal of Service Research* (13:3), pp. 267–282.
- Lin, J., Keogh, E., Fu, A., and Van Herie, H. 2005. "Approximations to Magic: Finding Unusual Medical Time Series," in *Proceedings of the 18th International Symposium on Computer-Based Medical Systems*, Washington DC, USA, pp. 329–334.
- Liu, B., and Zhang, L. 2012. "A survey of opinion mining and sentiment analysis," in *Mining Text Data*, C. C. Aggarwal and C. Zhai (eds.), New York: Springer, pp. 415–463.
- Liu, Y. 2006. "Word of Mouth for Movies Its Dynamics and Impact on Box Office Revenue," *Journal of Marketing* (70:3), pp. 74–89.
- Lotan, G. 2012. *KONY2012: See How Invisible Networks Helped a Campaign Capture the World's Attention*. <http://blog.socialflow.com/post/7120244932/data-viz-kony2012-see-how-invisible-networks-helped-a-campaign-capture-the-worlds-attention>. Accessed 25 October 2012.
- Ludwig, S., Ruyter, K. de, Friedman, M., Brüggem, E. C., Wetzels, M., and Pfann, G. 2013. "More Than Words: The Influence of Affective Content and Linguistic Style Matches in Online Reviews on Conversion Rates," *Journal of Marketing* (77:1), pp. 87–103.
- Malmgren, R. D., Hofman, J. M., Amaral, L. A. N., and Watts, D. J. 2009. "Characterizing individual communication patterns," in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Paris, France, pp. 607–616.
- Mislove, A., Marcon, M., Gummadi, K. P., Druschel, P., and Bhattacharjee, B. 2007. "Measurement and Analysis of Online Social Networks," in *Proceedings of the 2007 ACM SIGCOMM Internet Measurement Conference*, San Diego, USA.
- Mochalova, A., and Nanopoulos, A. 2013. "On The Role Of Centrality In Information Diffusion In Social Networks," in *Proceedings of the 21st European Conference on Information Systems*, Utrecht, Netherlands, paper 101.
- Mochalova, A., and Nanopoulos, A. 2014. "Restricting the Spread of Firestorms in Social Networks," in *Proceedings of the 22nd European Conference on Information Systems*, Tel Aviv, Israel, track 21, paper 2.
- Moe, W. W., and Trusov, M. 2011. "The Value of Social Dynamics in Online Product Ratings Forums," *Journal of Marketing Research* (48:3), pp. 444–456.
- Moon, S., Bergey, P. K., and Iacobucci, D. 2010. "Dynamic Effects among Movie Ratings, Movie Revenues, and Viewer Satisfaction," *Journal of Marketing* (74:1), pp. 108–121.
- Munzel, A., Jahn, B., and Kunz, W. H. 2012. "The Power of Saying Sorry – Insights on Customer Service in New Media Online Channels," in *Proceedings of the 12th International Research Conference in Service Management*, La Londe les Maures, France.
- Newman, M. E. J. 2002. "Spread of epidemic disease on networks," *Physical Review E* (66:1).

- Noufaily, A., Enki, D. G., Farrington, P., Garthwaite, P., Andrews, N., and Charlett, A. 2013. "An Improved Algorithm for Outbreak Detection in Multiple Surveillance Systems," *Statistics in medicine* (32:7), pp. 1206–1222.
- Pang, B., and Lee, L. 2008. "Opinion Mining and Sentiment Analysis," *Foundations and Trends in Information Retrieval* (2:1–2), pp. 1–135.
- Papazoglou, M. P., and Van Den Heuvel, W.-J. 2007. "Service oriented architectures: approaches, technologies and research issues," *The VLDB journal* (16:3), pp. 389–415.
- Peffers, K., Tuunanen, T., Rothenberger, M. A., and Chatterjee, S. 2007. "A Design Science Research Methodology for Information Systems Research," *Journal of Management Information Systems* (24:3), pp. 45–77.
- Pfeffer, J., Zorbach, T., and Carley, K. M. 2014. "Understanding online firestorms Negative word-of-mouth dynamics in social media networks," *Journal of Marketing Communications* (20:1/2), pp. 117–128.
- Poynter, R. 2008. "Facebook: The Future of Networking with Customers," *International Journal of Market Research* (50:1), pp. 11–12.
- Probst, F. 2011. "Predicting Users' Future Level of Communication Activity in Online Social Networks: A First Step towards More Advertising Effectiveness," in *Proceedings of the 17th Americas Conference on Information Systems*, Detroit, USA, paper 39.
- Probst, F., Grosswiele, L., and Pflieger, R. 2013. "Who will lead and who will follow: Identifying Influential Users in Online Social Networks," *Business & Information Systems Engineering* (5:3), pp. 179–193.
- Richardson, M., and Domingos, P. 2002. "Mining Knowledge-Sharing Sites for Viral Marketing," in *Proceedings of the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Edmonton, Canada.
- Rishika, R., Kumar, A., Janakiraman, R., and Bezawada, R. 2013. "The Effect of Customers' Social Media Participation on Customer Visit Frequency and Profitability An Empirical Investigation," *Information Systems Research* (24:1), pp. 108–127.
- Rozin, P., and Royzman, E. B. 2001. "Negativity bias, negativity dominance, and contagion," *Personality and social psychology review* (5:4), pp. 296–320.
- Saif, H., Fernández, M., He, Y., and Alani, H. 2013. "Evaluation datasets for Twitter sentiment analysis: a survey and a new dataset, the STS-Gold," in *Proceedings of the 1st International Workshop on Emotion and Sentiment in Social and Expressive Media*, Turin, Italy.
- Schmittlein, D. C., Morrison, D. G., and Colombo, R. 1987. "Counting Your Customers: Who-Are They and What Will They Do Next?" *Management Science* (33:1), pp. 1–24.
- Scholz, M., Dorner, V., Landherr, A., and Probst, F. 2013. "Awareness, Interest, and Final Decision: The Effects of User- and Marketer-Generated Content on Consumers' Purchase Decisions," in *Proceedings of the 34th International Conference on Information Systems*, Milan, Italy, paper 23.
- Sonnier, G. P., McAlister, L., and Rutz, O. J. 2011. "A Dynamic Model of the Effect of Online Communications on Firm Sales," *Marketing Science* (30:4), pp. 702–716.
- Stich, L., Golla, G., and Nanopoulos, A. 2014. "Modelling the spread of negative word-of-mouth in online social networks," *Journal of Decision Systems* (23:2), pp. 203–221.
- Stratus Contact Solutions 2014. *So What's next?* <http://stratuscontactsolutions.com>. Accessed 25 February 2014.
- Thet, T. T., Jin-Cheon Na, and Khoo, Christopher S. G. 2010. "Aspect-based sentiment analysis of movie reviews on discussion boards," *Journal of Information Science* (36:6), pp. 823–848.
- Tirunillai, S., and Tellis, G. J. 2012. "Does Chatter Really Matter? Dynamics of User-Generated Content and Stock Performance," *Marketing Science* (31:2), pp. 198–215.
- Venkatesh, V., Morris, M. G., Davis, G. B., and Davis, F. D. 2003. "User acceptance of information technology: Toward a unified view," *MIS Quarterly* (27:3), pp. 425–478.
- Von Neumann, J., and Morgenstern, O. 1947. *Theory of games and economic behavior*, Princeton, Ill.: Univ. Press.
- Wagner, M. M., Tsui, F.-C., Espino, J. U., Dato, V. M., Sitting, D. F., Caruana, R. A., McGinnis, L. F., Deerfield, D. W., Druzdzal, M. J., and Fridsma, D. B. 2001. "The emerging science of very early detection of disease outbreaks," *Journal of Public Health Management and Practice* (7:6), pp. 51–59.
- Wilson, C., Boe, B., Sala, A., Puttaswamy, K. P., and Zhao, B. Y. 2009. "User Interactions in Social Networks and their Implications," in *Proceedings of the Fourth EuroSys Conference*, Nuremberg, Germany.
- Wong, W.-K., and Cooper, G. 2003. "Bayesian Network Anomaly Pattern Detection for Disease Outbreaks," in *Proceedings of the 20th International Conference on Machine Learning*, Washington DC, USA, pp. 808–815.