

2nd International Conference on Information Technology and Quantitative Management,
ITQM 2014

Leveraging Localized Social Media Insights for Early Warning Systems

J. Bernabé-Moreno^a, A. Tejeda-Lorente^a, C. Porcel^b, E. Herrera-Viedma^{a,*}

^aDepartment of Computer Science and A.I., University of Granada, Granada E-18071, Spain

^bDepartment of Computer Science, University of Jaén, Jaén E-23071, Spain

Abstract

Social Media is more and more positioning as the channel to implement early warning systems for prompt reaction to events that might affect a company's core business. The enabling of geo-location for social media interactions unlocks new possibilities for decision making scenarios aiming at triage situation where the quality of service decays. In this work we provide a system and a set of soft metrics to quantify the impact created by the reaction of people directly affected by an incident in a particular area in order to facilitate the service providers' appropriate reaction, the decision making in marketing activities and to unveil customer acquisition opportunities applying the system to the competitors' customers.

© 2014 Published by Elsevier B.V. Open access under [CC BY-NC-ND license](#).

Selection and peer-review under responsibility of the Organizing Committee of ITQM 2014.

Keywords: Early-warning systems; Localized Social Media; Social Media Sensor; Social Media Insights

1. Introduction

Early warning systems have been successfully implemented in the last decades to support disaster prevention in countless scenarios, like earthquakes^{1,2}, pandemic expansion^{3,4}, flood and other natural hazards⁵, etc. The industry also discovered the competitive advantage resulting from implementing early warning systems, especially in the financial domain, where all variety of economic indicators have been used at a macro level to assess the vulnerability of emerging markets⁶ and to detect financial crisis in their early stages^{7,8}, but also at a much more micro level to detect for example critical transactions⁹, etc. Most of the companies providing services on a recurrent basis have the need for monitoring the service quality across the geographies where their customer base is located. Understanding the impact of a service quality decay on the set of affected customers, determining which reaction suits best to each incident, putting a reaction plan in place and executing upon the plan is key to prevent customers churn. Therefore, early warning systems have become an integral part of the service providers operations to implement customer retention strategies¹⁰.

* Corresponding author. Tel.: +34-958-24-4258 ; fax: +34-958-24-3317 .

E-mail address: viedma@decsai.ugr.es

The broad adoption of Social Media (SM) opens up new insights to feed into early warning systems. Customers increasingly engage with brands over the SM channels, and more and more service providers make use of SM channels to provide customer care or to almost real time monitor the reaction to a new product or campaign.

As network technologies started to support mobile data transfer, the interactions with the SM platforms became more pervasive and delocalized: everybody could post and consume SM content anytime from everywhere using a mobile device, such as a smartphone or a tablet. SM platforms rapidly provided the possibility of enhancing each piece of content created with a pair of coordinates, adding to the traditional time-stamp a *place-stamp*. Companies, that have been analyzing the SM channels at a rather broader geographical level for different use cases –e.g.: to understand what their customers and prospects think, where they see pain points, which campaigns are going well or which ones are under performing, etc–, can now have access to a completely new set of actionable insights at a much finer regional scale.

Understanding where something has been posted and analyzing all content created in a particular area is a valuable information source to feed into early warning systems of any kind. In this paper we define a set of metrics to quantify the impact of the SM interactions created in a particular place with respect to an entity –that can be a brand, a company, an institution, etc–, supporting also the separation by purpose (e.g.: complaints, criticism, information request, etc). Using such metrics, we designed a system to support decision making in the realm of marketing activities, acquisition of new customers and retention of existing ones. To prove our metrics, we built a monitoring system which computes the metrics for a places over time based on Twitter data. To illustrate how our metrics work, we provide a real-world example based the British Transportation System.

We have organized this paper starting with the background for our research presented in Section 2). In Section 3, where we define our whole set of metrics. After that, Section 4 deals with the system implementation and Section 5 presents the results obtained for a real world scenario. We closure the paper with the conclusions where further research lines are also pointed out.

2. Background

SM is being broadly adopted as a channel to get prompt and unfiltered feedback from a company's customers base and consequently, there is more and more literature about researches on different aspects of this adoption: in the seminal paper¹¹, apart from declaring the integration with the traditional media to reach customers as a must, Kaplan et al. point SM as the channel to engage with customers in a time-close and high-efficient manner. In¹² the authors declare SM as integral part of the promotion mix highlighting the less controlled and hence more insights revealing nature of customers interactions. In¹³ a thorough analysis on how all the interactions created over Twitter about a brand impacts its corporate image. Twitter, founded in March 2006, is the microblogging platform *par excellence*; users can send and read tweets or text messages containing maximum number of 140 characters. Optionally, users can also tweet pictures, videos and URLs, or re-tweet what other users tweeted. The adoption of SM appears to be a game changer when it comes down to engaging with customers: in¹⁴ the authors explained how microblogging was shaking traditional business models by increasing the role of product quality, as the time window where product new adopters didn't have any feedback on the product gets drastically reduced. In¹⁵ one interesting aspect based on communication intentions is presented to confirm that similar intentions foster connectivity between users.

As proved in¹⁶ and¹⁷, the spreading of bad news takes place really fast over the SM channel, which makes them of great value as a source to set early warning system upon for early detection of customers' complaints, service outages, etc. In¹⁸, Sakaki et al. define an algorithm based on particle filtering for geo-location and spread for earthquakes early detection based on tweets. Also based on tweets, Culotta et al. suggest in¹⁹ a method to detect epidemic expansion on early stages. In²⁰ a stochastic model for dynamic of the interactions based on the underlying network structure is employed to generate useful predictions about the spread of information.

Our Impact metric, relies on how influential a particular SM user is. Modeling influence in SM channels has been subject of intense research over the last few years. Kwak²¹ defined 3 metrics aimed at quantifying the *social influence*: the so called *propagation influence*, based on the Google Search PageRank algorithm²², *followers influence* –more followers implies more influence–, and *re-tweet influence* –more re-tweets means more influence–. Ye and Wu²³ relied on the same set of metrics but changing the propagation influence by a much simpler to compute *reply influence* –the more replies one user receives, the more influential the user is–. Cha²⁴ also identified 3 influence drivers: the size

of the user's audience or social network -*indegree influence*-, the generated content with pass-along value -*retweet influence*-, and the engagement in others' conversation -*mention influence*-. Romero et al.²⁵ develop a mechanism to quantify how the exposure to other users is making them adopt a new behavior.

3. Framework definition

The ultimate aim of our framework is providing a means to quantify the impact in an efficient way, so that our metrics can be consumed near real time by early warning systems for decision making. The Impact of a SM interaction with a brand can be modeled by reach -or number of exposed users to this interaction- and a so called *differential perception factor*, what has been introduced to remove the SM behavioral bias at user level typically present in the SM networks. To explain it in an intuitive manner, a complaint coming from a particular user who is always complaining is perceived as less critical than a complaint from somebody who hardly ever posts anything negative about any service. As a remarkable side effect, the impact contribution from potential spam users created to deliberately damage a brand image is minimized.

The Impact computed over all users located in a place provides a really sensible Key Performance Indicator (KPI) to take decisions upon. In our approach, the Impact is provided in different categories, which perfectly maps with the way big corporations are usually structured in departments. For example, the complaint management department is interested in monitoring the impact over time of the complaints coming from a place over the SM channel, whereas marketing rather focuses on the monitoring of suggestions, criticism and engagement with running campaigns.

3.1. Preliminary definitions

Before jumping into the framework definition, we are going to establish a set of concepts our metrics build upon:

Definition 1. The set U represents the set of Social Media Users from which we have evidence they have been in the location ($InLocation(u_i, \Delta t)$) we are monitoring during the time period under analysis Δt

$$U \equiv \{u\}, \forall u_i \in U, InLocation(u_i, \Delta t) \quad (1)$$

Definition 2. The Social Network for a given user u_i is defined as:

$$SN(u_i) \equiv \{u\}, \forall u_j \in SN(u_i), Follows(u_i, u_j) \quad (2)$$

$Follows(u_i, u_j)$ is a function representing a SM connection between the users u_i and u_j , so that u_i is exposed to the SM content generated by u_j . $Follows(u_i, u_j)$ is not always commutative; although in several SM platforms it is the case (e.g.: Facebook or Linked.in). There are others where it is not necessarily the case, like Twitter, where $Follows(u_i, u_j) \not\Rightarrow Follows(u_j, u_i)$

Definition 3. The set $SN(U)$ represents the set of all the users being followed by the users in U :

$$SN(U) \equiv \{u\}, \forall u_i \in SN(U), \exists u_j \in U | u_i \in SN(u_j) \quad (3)$$

Definition 4. We define all user interactions ($UserInteractions$) for a given user u_i over a time interval Δt , as:

$$UserInteractions(u_i, \Delta t) \equiv \{it\}, \forall it_i \in UserInteractions(u_i, \Delta t), Author(u_i, it_i, \Delta t) \quad (4)$$

A Social Media interaction represents the atomic piece of content generated by the user u_i during the time Δt in a Social Media Platform (e.g.: a tweet, a re-tweet). Thus, $Author(u_i, it_i, \Delta t)$ is a function that retrieves *True* if u_i created the interaction it_i in the time period Δt , and *False* otherwise. The time interval t might be measured in weeks, days or hours, depending on the use case and consists of two extremes: $t_startdate$ and end date $t_enddate$.

Definition 5. A Social Media Entity E is the representation of the set of all terms used by Social Media Users to interact with a real world entity such as a brand, a corporation, an institution, a club, etc. It includes for example social media account name(s), product names, company abbreviations or company slogans.

Definition 6. We define the set of Interactions for a given user u_i with the entity E over a time interval Δt as:

$$Interactions(u_i, E, \Delta t) \equiv \{it\}, \forall it_i \in Interactions(u_i, E, \Delta t), Author(u_i, it_i, \Delta t) \wedge related(it_i, E) \quad (5)$$

Where $related(it_i, E)$ is a NLP membership function retrieving *True* if the iteration it_i is connected to the entity E -intuitively, one or more words from the Entity defining set are mentioned in it_i - and *False* otherwise.

3.2. User-Entity engagement

Based on the before mentioned definitions, we introduce the concept of “engaged”, defined as a logical function:

$$\text{Engaged}(u_i, E, \Delta t) \equiv \text{True}, \exists it_i, it_i \in \text{Interactions}(u_i, E, \Delta t), u_i \in U \cup SN(U) \quad (6)$$

Where u_i is the user, E is the representation of the Entity, Δt is the time span specified consisting of two components ($t_{startdate}$ and $t_{enddate}$), it_i represents a social media interaction and $\text{Interactions}(u_i, E, \Delta t)$ represents the interactions of the user u_i related to the Entity E in the time interval Δt , as we explained before. At user level, it's also possible to define a metric to quantify the level of engagement of the user with the Entity:

$$\text{EntityEngagementIndex}(u_i, E, \Delta t) = \frac{\# \text{Interactions}(u_i, E, \Delta t)}{\#(\cup_{k=1}^{\#E} \text{Interactions}(u_i, E_k, \Delta t))} \quad (7)$$

Where u_i represents a given SM user, E is the representation of the Entity, $\text{Interactions}(u_i, E, \Delta t)$ is as defined before and $\#(\cup_{k=1}^{\#E} \text{Interactions}(u_i, E_k, \Delta t))$ is the cardinal for the union set of all interactions with all possible entities created by the user u_i during the time span Δt

The Entity Engagement Index can be approximated by:

$$\text{EntityEngagementIndex}(u_i, E, \Delta t) = \frac{\# \text{Interactions}(u_i, E, \Delta t)}{\# \text{Interactions}(u_i, \Delta t)}$$

3.3. Social Media Communication Intent

Each and every SM interaction resulting in the creation and diffusion of content has an underlying purpose: praise a piece of information or a company or an action, express some criticism, make a direct complaint, request information, provide an answer, etc. In the same way we introduced before the concept of *Social Media Entity*, we now provide the definition for *Communication Purpose Category*

Definition 7. A *Communication Purpose Category* P is the representation of the set of all terms in all varieties of forms used by Social Media Users to express the intent represented by the Category.

Even if the boundaries might not be crisp, we can assign each interaction to a *leading Purpose Category*:

$$\forall it_i \in \text{Interactions}(u_i, E, \Delta t), \exists p_k, \text{Purpose}(it_i) = p_k, p_k \in PC \quad (8)$$

Where it_i represents a SM interaction, $\text{Interactions}(u_i, E, \Delta t)$ is the set of all interactions created by u_i over Δt , p_k is a the leading Communication Purpose, PC is the set of all Communication Purpose Categories.

3.4. Early Warning Metrics

Based on the concepts introduced in the previous sections 3.2 and 3.3, we can define a set of metrics to quantify the impact created by the users located in a given area over time, and thereby enable the early reaction and steering.

We introduce the so called *Differential Perception Factor* modeled as *Purpose Share* which allows for latterly defining a correction factor to remove the SM behavioral bias:

$$\text{DPF}(u_i, E, P, \Delta t) = \frac{\#(\text{Interactions}(u_i, E, \Delta t) \cap \text{Interactions}(u_i, P, \Delta t))}{\# \text{Interactions}(u_i, P, \Delta t)} \quad (9)$$

To make it more intuitive, let's bring up one example: let's assume that a given user in a location started posting complaints over Twitter about the bad services provided by his/her mobile operator. If the same user was very active posting complaints about many other companies such as the local transportation service, the internet provider, the employer, certain celebrities, etc, the Purpose Share for *Complaints* would be rather low. On the other hand, if the same user hardly ever complaints about anything, a single interaction pointing out his/her discontent with the mobile operator would be perceived as something rather serious and more significant.

The impact measure of a social media interaction originated in a particular area shall consider the number of users that are exposed to this content, no matter if they are in the same area or some where else.

$Exposed(u_i, u_j, E, \Delta t)$ is a logical function defined as:

$$Exposed(u_i, u_j, E, \Delta t) = \begin{cases} True, & u_j \in SN(u_i), \exists it_k, it_k \in Interactions(u_i, E, \Delta t), \\ & P(read(u_j, it_k, \Delta t)) \geq Threshold \\ False, & otherwise \end{cases} \quad (10)$$

where $P(read(u_j, it_k, \Delta t))$ is the probability that the user u_j reads the content posted in the interaction it_k in the designated time Δt . The $Threshold \in [0, 1]$ is defined to narrow down the selection.

The reason why we introduce the concept of *Exposed User* is to address the fact that not all the SM content created by the social network of a particular user is consumed by the user. The subset of users exposed to the topic can then be defined as:

$$ExposedUsers(u_i, E, \Delta t) \equiv \{u\}, \forall u_j, Exposed(u_i, u_j, E, \Delta t) = True, u_i \in U \quad (11)$$

Based on the DPF and on the number of people exposed to the SM interaction, we can define Impact as:

$$Impact(u_i, E, P, \Delta t) = \mathfrak{I}(EntityEngagementIndex(u_i, E, \Delta t), DPF(u_i, E, P, \Delta t), \#ExposedUsers(u_i, E, \Delta t)) \quad (12)$$

The \mathfrak{I} function is usually a simple product but can also be implemented in a more sophisticated way giving for instance different weights to the components.

As in certain scenarios is more critical having very quickly a probably *not-that-precise* value to act upon, than a high-precision metric but also with higher latency, there are some approximations that can be done. When the trade-off between precision and time-to-results is decided for the second, the DPF can be simplified as:

$$DPF(u_i, E, P, \Delta t) \approx 1 \quad (13)$$

Obviously, the ultimate purpose of DPF to remove the SM behavioral bias is thereby annulled. One of the most time consuming processes is the computing of the Exposed Users set. As what it's really required in the Impact function as we defined before is the cardinal of the set, an approximation using a correction coefficient on the number of users that are part of $SN(u_i)$ removes the complexity derived from computing the probability:

$$\#ExposedUsers(u_i, \Delta t) \approx \#SN(u_i) * K, K \in [0, 1] \quad (14)$$

The resulting metric to take action on is defined as an aggregation over the individual Impact resulting into a quite big number:

$$Impact(U, E, P, \Delta t) = \sum_{i=1}^{\#U} Impact(u_i, E, P, \Delta t) \quad (15)$$

In order to make this Impact metric more actionable, the value can be mapped to categories, applying different Levels –defined by a pair of min and max value– (e.g.: the typical (R)ed, (A)mber, (G)reen scala, etc). To select the category boundary values is advisable to have a long enough history available to understand how the values change over time. Even if one could define the category boundary values generically for all the places to be monitored, the heterogeneity among geographical areas might introduce the need for location-specific RAG boundaries definition.

4. System Description

Before running the system, a set of configuration parameters needs to be supplied, such as the categorization of the Entity to be monitored, the places to inspect, the set of brand-specific or industry-specific purpose categories semantic fields and the time unit for the insights aggregation.

The end-to-end process consists of several steps with clearly defined purpose:

1. Content Polling: extracts from the SM platform the content generated in the place(s) under monitoring and stores it for further processing

2. **Content Tagging:** flags the interactions that are related to the entity we are interested in and assigns a Communicative Purpose Category to them
3. **Users Information Polling:** gathers all relevant information about the users authoring the interactions and their SM networks
4. **Metrics Computing:** applies the set of metrics defined in the previous section to obtain the impact values and eventually provides the mapping to the categories.

In Figure 1 we show the modules of the system based on the previous metrics: *Content Harvesting*, *Content Tagging*, *User Info Gathering*, *Metrics Creation*. In the following subsections we are going to describe each and every step, explaining which modules are required and providing details about the implementation.




Fig. 1: System structure

4.1. Content Harvesting

Relying on the Twitter Search API publicly available Twitter Search API¹ the harvester picks all tweets created in a given area, which is defined as a pair of geo-coordinates and a radius as part of the system configuration. A pre-filtering by language can also be applied to the harvester to just pick tweets in a given language.

4.2. Content Tagging

Once all the tweets created in an area have been gathered, the tagging module separates all tweets related to the brand under monitoring from the rest. The separation relies on finding occurrences in the interaction content of terms supplied in the brand definition file. These terms are account names, hashed tags employed to identify the brand in social media channels, etc (e.g.: if we consider the German airline Lufthansa, @lufthansa, all regional accounts associated to Lufthansa like @Lufthansa_DE, @Lufthansa_BR, @Lufthansa_AR, hashtags like #lufthansa and the name of the services they are offering, in this case, the flight codes LH6670, LH6671, etc as well as the programs run by the company, like @Miles_and_More, @MilesandMore and #milesandmore)

In order to provide certain tolerance when the users enter the name of the brand, the tagging module works with a string similarity function²⁶ to accept spelling mistakes (e.g.: *lufthansa* or *lufhansa* with a similarity over 0.6 wouldn't be rejected if the threshold was set to 0.6)

Tweets tagged positively as related to the brand are assigned a communication purpose category applying the same technique. The definition of the categories is up to the use case to be implemented on top of the generated metrics. Thus, categories like praise, criticism, information requests, suggestions, etc make only sense if the brand has specialized departments at its disposal to handle. In the simplest case, a mere sentiment-like categorization separating positives from negatives could be helpful. The category definition file is not a trivial task due to the underlying complexity in the Natural Language Processing task. A good strategy is the n-gram extraction^{27,28} over a long history of content related to a particular category, ideally followed by a supervised step (e.g.: forums are best suitable for the extraction and are usually divided into threads, that very well map to purpose categories). Disambiguation is handled

¹ Available at <https://dev.twitter.com/docs/api/1/get/search?>

relying on both Part of Speech tagging and the presence of more than one terms related to the Entity or Purpose Category.

4.3. User Info Gathering

Our impact metric is an aggregation of the individual impact generated by each user who has authored one of the posts flagged as *related to the brand* in the previous step.

The *User Information Gathering* module consults the SM Platform API to retrieve the meta information required at user level, including their social network.

If the approximations suggested in the equations 13 and 14 are considered viable for the use case, this module is configured to just gather the necessary information, resulting in a much better performance but trading off certain precision. Especially the process of determining whether a given user $u_j \in SN(u_i)$ belongs to the set of $ExposedUsers(u_i, \Delta t)$ is particularly time consuming. We implement it by defining a time window centered on a SM interaction (e.g.: 120 minutes) and then checking whether there is a SM interaction $it_j \in Interactions(u_j, \Delta t)$ user u_j , whose time window $[t(it_j - 60min), t(it_j + 60min)]$ overlaps with the time window of any of the interactions created by u_i , $\exists it_i \in Interactions(u_i, E, \Delta t), [t(it_j - 60min), t(it_j + 60min)] \cap [t(it_i - 60min), t(it_i + 60min)] \neq \emptyset$. Obviously, it requires gathering all the transactions from the user u_i and from all other users in the $SN(u_j)$ during the period of time Δt and computing the overlapping, which might compromise the performance of the system.

4.4. Metrics Computation

With all the relevant interactions available and properly tagged by Entity and by Communication Purpose, as well as the information required about the authors of these transactions and their SM network, the module in charge of creating the metrics can proceed: for each author involved in a interaction flagged as relevant as explained in the previous section 4.2, the Impact according to the equation (12) is computed. It requires the previous calculation of the single components: the Entity Engagement Index (equation 7), the DPF (equation 9) and the size of the Exposed Users (equation 11).

Once the individual Impact has been computed, the overall Impact is obtained applying the aggregation (eq. 15).

This module can also map the value obtained to one of the impact categories –whenever available– to make the resulting number more actionable (as explained in the section 3.4)

5. Evaluation Results

In order to prove our results, we set up 2 harvesters in the two main airports in the city of London: Heathrow and Gatwick (centered on the airports with a radius of 5 km). The harvesters gathered between the 23th of November 2013 and the 23rd of January 2014 a total of 852319 SM interactions.

We have chosen several Entities within the same sector, namely railway transportation, mainly because of two reasons: the amount of people using trains on a regular basis is significantly large and the customer satisfaction is usually low, which push people to express their discontent over the SM channels. We considered Virgin Trains, First Capital Connect, National Rail, the companies offering exclusive express services Gatwick Express and Heathrow Express, and the local operator Southern

As Communication Purpose Category we selected *Complaints*, as mentioned before. The semantic field required for classifying interactions by purpose for the category *Complaints* has been pulled with a n-grams extraction based semi-automatic by frequency from forums and SM content from the 6 before mentioned company Twitter accounts. The classification is also supported by a basic natural language processing to remove the stopwords, tokenize and lematize on top of the extracted n-grams, etc –the particular NLP-related details remain outside the scope of this paper–. In the Figure 2 we can see the top 20 words based on their penetration over all SM interactions related to the before mentioned entities flagged positively as *complaints*. Similarly, the Entities have been modeled including all relevant account information, hashed tags and even non-official accounts, like @SouthernTrains created as a parody of the official @SouthernRailUK. In Figure 3 we provide for both harvesters, the amount of interactions assigned to the corresponding brands (first row) and the subset of those classified as a complaint. The numbers reflect the reality

<i>Top 10</i>	delay	no train	taxi	cancel	shut	closed	disrupt	flood	broken	late	problem
	10,48%	7,16%	4,98%	4,87%	3,75%	3,09%	2,69%	2,63%	2,05%	2,00%	1,97%
<i>Top 20</i>	break	fire	miss	f**k	affect	alter	fault	bad	chaos	stuck	shit
	1,92%	1,84%	1,82%	1,62%	1,54%	1,47%	0,95%	0,92%	0,86%	0,79%	0,74%

Fig. 2: Penetration of the top 20 words categorizing a complaint across carriers

of the railway transportation for both airports. Heathrow is only connected to London by Tube –not included in this analysis– and by the exclusive Heathrow Express Service. The Gatwick Airport Railway Station is an important node in the British railway infra-structure offering long-distance trains (Southern), First Capital Connect trains, the Gatwick Express to London Victoria, etc. The numbers confirm Gatwick as a much heavier station.

	All	Virgin Trains	Southern Rail	National Rail	FirstCC	Gatwick Express	Heathrow Express
Gatwick Airport							
Total	10983	12	1185	8717	565	504	0
Complaints	5911	4	545	4779	352	231	0
Heathrow Airport							
Total	2817	26	17	2430	9	0	335
Complaints	668	8	10	572	2	0	75

Fig. 3: Number of Interactions per Twitter Harvester in Total and identified as a Complaint

Figures 4 and 5 shows the Impact metric computed for all 6 entities over 2 months in both airports. The highest value (over 600K) originated on Dec 24, Entity: National Railway, reflects the train service disruption when the storm Emily was striking the country ². The second highest (ca. 400K) registered also for National Railway on Jan 17 is due again to weather causing flooding ³. Obviously extreme service disruptions lead to the corresponding reaction in the media, which gets reflected in our metric, but we can now quantify the impact over time and compare the impact of reaction to different events in different days (e.g.: the storm had a much bigger impact than the flooding, as we can see). According to our charts, each single Entity has been heavily impacted by the storm, but the Gatwick Express Services. The reason is because it remains closed between Christmas and New Year ⁴. Our metric can also quantify the impact of small decays in the quality of service. Figure 6 shows the total delay in minutes accumulated day by day by the First Capital Connect train lines to or over Gatwick airport. The days with high delay values are usually reflected as peaks in the Impact curve for FCC shown in Figure 5 –Nov 26, Nov 29, Dec 5, etc–, yet the Impact intensity does not necessarily correlate with the delay in minutes or with the number of cancellations. This is where we prove how valuable our metrics are: in addition to the hard KPI –like minutes of delay in taking Gatwick as a reference point–, we can feed early warning systems for decision making systems with a soft KPI which quantifies the Impact the delays in Gatwick are having on the brand image in the social media channels.

6. Conclusions

In this work we suggest a metric to quantify the impact that a localized customers community of a service provided by a company, have on the company's image. To build this metric, we take into account all SM interactions created in a particular area related to the company or the service provided by the company, the underlying communication purpose per interaction and how the authors of these interactions are connected to other SM users.

² Southern services suspended due to strong storm on December 24th <http://www.bbc.com/news/uk-25785804>

³ News about the rail services disruption in south-east England due to heavy flooding on January 7th <http://www.bbc.com/news/uk-25785804>

⁴ <http://www.londontoolkit.com/blog/daytrips/london-tours-on-christmas-day-new-years-2012/>

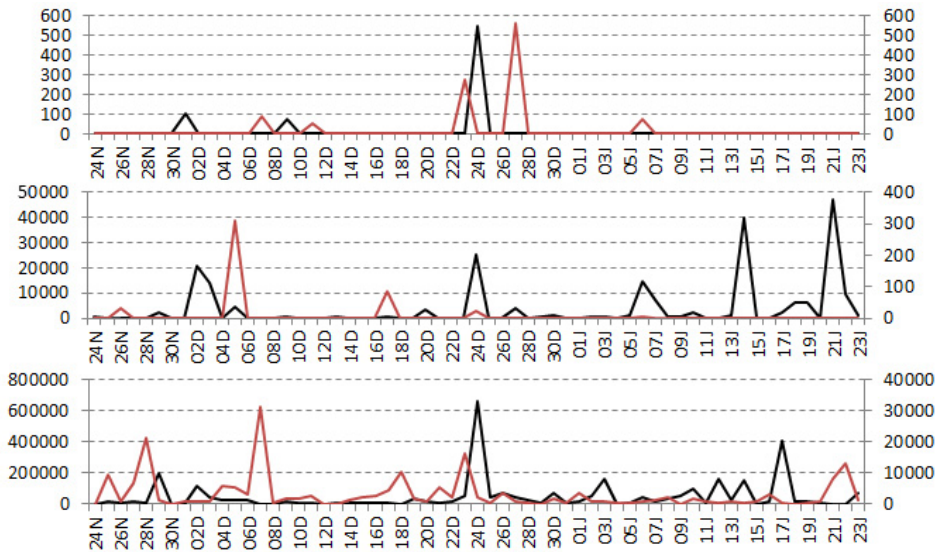


Fig. 4: Impact over 2 months for Virgin Trains (top), Southern Railway (mid) and National Rail (bottom) for the places Heathrow Airport -red- and Gatwick Airport -black-




Fig. 5: Impact over 2 months for Gatwick Express (top), FCC (mid) and Heathrow Express (bottom) for the places Heathrow Airport -red- and Gatwick Airport -black-




Fig. 6: Delay and Cancellations for First Capital Connect trains via Gatwick over 2 months. Source: <http://www.firstcapitalconnect.co.uk>

Additionally, we suggest the mapping of the metric value to a category or level to make it actionable and ready to be fed into an early warning system. We also address the cases where the time to results is critical by providing approximations to the single components of our metric and removing therefor the time consuming steps but trading some precision off.

Our approach treats each interaction the same way independently on the content. Establishing a classification of the criticality based on the interaction message or including a new dimension based on sentiment analysis to the Impact definition are two meaningful research lines for future work. Whereas we focused on early reporting based on metrics monitoring, further research may cover predictive modeling on top and broaden therefore the application domain.

Acknowledgements

This paper has been developed with the financing of projects TIN2010-17876, TIC5299 y TIC-5991.

References

- Allen RM, Kanamori H. The potential for earthquake early warning in southern california. *Science* 2003;**300**(5620).
- Wu YM, Teng TI. A virtual subnetwork approach to earthquake early warning. *Bulletin of the Seismological Society of America* 2002;**92**(5).
- Bouma MJ, Kaay HJ. The ei niño southern oscillation and the historic malaria epidemics on the indian subcontinent and sri lanka: an early warning system for future epidemics? *Tropical Medicine & International Health* 1996;**1**(1).
- Thomson M, Doblas-Reyes F, Mason S, Hagedorn R, Connor S, Phindela T, et al. Malaria early warnings based on seasonal climate forecasts from multi-model ensembles. *Nature* 2006;**439**(7076).
- Basher R. Global early warning systems for natural hazards: systematic and people-centred. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 2006;**364**(1845).
- Goldstein M, Kaminsky GL, Reinhart CM. *Assessing financial vulnerability: an early warning system for emerging markets*. Peterson Institute; 2000.
- Bussière M, Fratzscher M. Towards a new early warning system of financial crises. *Journal of International Money and Finance* 2006;**25**(6).
- Edison HJ. Do indicators of financial crises work? an evaluation of an early warning system. *An Evaluation of an Early Warning System (July 2000) FRB International Finance Discussion Paper* 2000;(675).
- Scheffer M, Bascompte J, Brock WA, Brovkin V, Carpenter SR, Dakos V, et al. Early-warning signals for critical transitions. *Nature* 2009; **461**(7260).
- Reichheld FP, Sasser WE. Zero defections: Quality comes to services 1990;.
- Kaplan AM, Haenlein M. Users of the world, unite! the challenges and opportunities of social media. *Business horizons* 2010;**53**(1).
- Mangold WG, Faulds DJ. Social media: The new hybrid element of the promotion mix. *Business horizons* 2009;**52**(4).
- Jansen BJ, Zhang M, Sobel K, Chowdury A. Twitter power: Tweets as electronic word of mouth. *Journal of the American society for information science and technology* 2009;**60**(11).
- Hennig-Thurau T, Wiertz C, Feldhaus F. Exploring the twitter effect: an investigation of the impact of microblogging word of mouth on consumers early adoption of new products. *Available at SSRN 2016548* 2012;.
- Java A, Song X, Finin T, Tseng B. Why we twitter: understanding microblogging usage and communities. In: *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis*. ACM; 2007. .
- Naveed N, Gottron T, Kunegis J, Alhadi AC. Bad news travel fast: A content-based analysis of interestingness on twitter. In: *Proceedings of the 3rd International Web Science Conference*. ACM; 2011. .
- Park J, Cha M, Kim H, Jeong J. Managing bad news in social media: A case study on domino's pizza crisis. In: *ICWSM*. 2012. .
- Sakaki T, Okazaki M, Matsuo Y. Earthquake shakes twitter users: real-time event detection by social sensors. In: *Proceedings of the 19th international conference on World wide web*. ACM; 2010. .
- Culotta A. Towards detecting influenza epidemics by analyzing twitter messages. In: *Proceedings of the first workshop on social media analytics*. ACM; 2010. .
- Colbaugh R, Glass K. Early warning analysis for social diffusion events. *Security Informatics* 2012;**1**(1).
- Kwak H, Lee C, Park H, Moon S. What is twitter, a social network or a news media? In: *Proceedings of the 19th International Conference on World Wide Web; WWW '10*. New York, NY, USA: ACM. ISBN 978-1-60558-799-8; 2010. .
- Page L, Brin S, Motwani R, Winograd T. The pagerank citation ranking: bringing order to the web. 1999;.
- Ye S, Wu SF. Measuring message propagation and social influence on twitter.com. In: *Proceedings of the Second International Conference on Social Informatics; SocInfo'10*. Berlin, Heidelberg: Springer-Verlag. ISBN 3-642-16566-4, 978-3-642-16566-5; 2010. .
- Cha M, Haddadi H, Benevenuto F, Gummadi K. Measuring user influence in twitter: The million follower fallacy. In: *4th International AAAI Conference on Weblogs and Social Media (ICWSM)*. 2010. .
- Romero DM, Meeder B, Kleinberg J. Differences in the mechanics of information diffusion across topics: Idioms, political hashtags, and complex contagion on twitter. In: *Proceedings of the 20th International Conference on World Wide Web*. New York, NY, USA: ACM. ISBN 978-1-4503-0632-4; 2011. .
- Yang QX, Yuan SS, Zhao L, Chun L, Peng S. Faster algorithm of string comparison. *Pattern Analysis & Applications* 2003;**6**(2).
- Cavna WB, Trenkle JM, et al. N-gram-based text categorization. *Ann Arbor MI* 1994;**48**113(2).
- Fürnkranz J. A study using n-gram features for text categorization. *Austrian Research Institute for Artificial Intelligence* 1998;**3**(1998).