Universidad de Granada

Departamento de Ciencias de la Computación
e Inteligencia Artificial

Programa de Doctorado en Ciencias de la Computación
y Tecnología Informática

New methods for Knowledge Discovery in Geo-localized Social Media Networks

Tesis Doctoral

Juan Bernabé Moreno

Granada, Mayo 2015

Editorial: Universidad de Granada. Tesis Doctorales
Autor: Juan Antonio Bernabé Moreno
ISBN: 978-84-9125-238-2
URI: http://hdl.handle.net/10481/40800

UNIVERSIDAD DE GRANADA

New methods for Knowledge Discovery in Geo-localized Social Media Networks

MEMORIA PRESENTADA POR

Juan Bernabé Moreno

PARA OPTAR AL GRADO DE DOCTOR EN INFORMÁTICA

Junio de 2015

DIRECTORES

Enrique Herrera Viedma y Carlos Porcel Gallego

Departamento de Ciencias de la Computación
e Inteligencia Artificial

La memoria titulada "*New methods for Knowledge Discovery in Geo-localized Social Media Networks*", que presenta D. Juan Bernabé Moreno para optar al grado de doctor, ha sido realizada dentro del Programa Oficial de Doctorado en "*Ciencias de la Computación y Tecnología Informática*", en el Departamento de Ciencias de la Computación e Inteligencia Artificial de la Universidad de Granada bajo la dirección de los doctores D. Enrique Herrera Viedma y D. Carlos Porcel Gallego.

El doctorando y los directores de la tesis garantizamos, al firmar esta tesis doctoral, que el trabajo ha sido realizado por el doctorando bajo la dirección de los directores de la tesis, y hasta donde nuestro conocimiento alcanza, en la realización del trabajo se han respetado los derechos de otros autores a ser citados cuando se han utilizado sus resultados o publicaciones.

Granada, Junio de 2015

El Doctorando Los directores

Fdo: Juan Bernabé Moreno Fdo: Enrique Herrera Viedma Fdo: Carlos Porcel Gallego

Esta tesis doctoral ha sido ha sido subvencionada por los proyectos TIC-5991 and TIN2013-40658-P.

Dedicada a mi padre:
D. Juan Antonio Bernabé Llorente

Agradecimientos

Esta tesis está especialmente dedicada a mi padre, D. Juan Antonio Bernabé Llorente, porque con ella conseguimos cerrar una asignatura pendiente de los dos. A mi madre y hermana, por su apoyo incondicional duranto todos estos años de períodos intermitentes e intensos.

En segundo lugar, mencionar el núcleo duro de esta tesis... Al gran Enrique, que a parte de ser altamente citado y darle prestigio a todos los papers en los que he tenido la fortuna de trabajar con él, es una persona irrepetible siempre interesado en abrir nuevas vías y explorar nuevos caminos. Al inigualable Carlos, que con su experiencia ha sabido siempre afinar todos los instrumentos que hemos puesto en sus manos. Y a mi AMIGO con mayúsculas Álvaro, compañero de fatigas y de cervezas en Munich y en Granada sin el que las veces que he caído no me hubiera podido levantar.

Quiero extender la dedicatoria a mis amigos Rubén y Rocío (y al pequeño Diego), por los momento de Data Science compartidos y porque a pesar de la distancia y el tiempo siempre han estado ahí.

Jetzt mache ich weiter mit Christoph, ein guter Freund, der mit mir viel durchgemacht hat und der immer die passende Antwort fuer alle Technolgiefragen parat hat. Bei Dir moechte ich mich ganz herzlich bedanken.

Of course I'm not going to forget my good friend and ex-boss Daryl... who always believed in the innovation power of my research and who helped me a lot all along this period. Thank you Daryl!

No puedo olvidar en este mensaje de agradecimiento a mis amigos: Zafra, Baldo, Manolo y Chicho, que sin haber aportado mucho en el terreno academico, se han hecho cargo de todos los demás aspectos durante la realizacin de esta tesis.

Quería dejar para el final, un agradecimiento muy especial, a la persona que ha "sufrido" esta tesis durante muchos años más que nadie. Sin su cariño, su compresión y su motivación durante todo este tiempo nunca hubiese sido capaz de hacerla. Gracias Kathrin!

GRACIAS A TODOS

Table of Contents

Chapter I

PhD dissertation

1. Introduction

It's difficult to think of today's digital world without the existence of Social Media (SM). SM is increasingly becoming an important part of our lives in a more and more integrative way. Twitter's mission statement captures in a sentence to which extend SM has revolutionized the way we create and consume information *"Our mission: To give everyone the power to create and share ideas and information instantly, without barriers"* [1].

SM is in a way a digital register of portions of everyone of us' lives, which put together, can be seen as a living document of our culture; even the Library of Congress is archiving all tweets sent by Americans ever since the Twitter platform went live. Lee Humphreys, a communication professor at Cornell University in New York, said that "the brief online messages -tweets- can reveal volumes about the culture where they were produced" [2].

The usage of SM is an ever-growing phenomenon [Whi13]. Media consumers have been increasingly shifting from classic (printed) media to digital platforms. As a result the communication stopped being one-way, with clearly defined *author / reader* roles. With the advent of the web 2.0, the definition of *author* started to blur [O'r07]. The blogosphere empowered readers to make their own contributions to the content published by a given author, which radically increased the information richness, adding further perspectives and points of view. Simultaneously, media started to be democratized, as anybody could start a blog and the visibility of the blog in the search engines was determined a priori by the number of people that considered the blog to be relevant outside the realm of the paid search [PBMW98].

The SM platforms based on the concept of *micro-blogging* took it to the next level, as everybody could be an author and a reader anytime. The *push-first, comment-later* paradigm so popular in the blogosphere started to look old-fashioned. Rather, anybody was empowered to initiate a communication, enrich an existing thread, jump from a thread to another one, ignore, criticize, share richer content like pictures, videos, etc.

The ease of publishing, sharing and consuming content boosted the adoption of these social media platforms as the place to talk any time about everything with everybody. The best example is Twitter, which has become a communication platform for almost all the digital world [KLPM10].

[1] https://about.twitter.com/company

[2] http://www.businessinsider.com/library-of-congress-is-archiving-all-of-americas-tweets-2013-1

By March 2012 the platform counted 140 million active users creating an average of 340 million tweets a day [Ben12]. The night of November 7th, during 8:11 and 9:11 pm when the world wanted to share the results of the US elections, an average of 9,965 Tweets per Second (TPS) [3] resulted in the creation of more than 35 million tweets within one hour.

What makes SM so revealing, unlike other channels, is the fact that users are less reluctant to express -almost in a unfiltered way- what is literally going through their minds [MZ12] –with the exception of work-oriented and other specific purpose platforms,such as LinkedIn–. And no matter how empty of content a message seems to be, each and every SM interaction encapsulates a piece of communication with an underlying intent. A SM interaction can be seen as the result of basically two choices made by the author: the what to convey and the how to transcribe it in words -or other richer assets, such as pictures, videos, etc-. Even the SM specific forms of interaction, such as the famous Facebook "like" or Twitter "favourite" or "retweet" are full of meaning.

In its essence, SM started as a space where anybody with an account could interact with any other user, share content, express their own personal views, etc. without being subjected to any kind of censorship. As a side effect of this democratization of the Web, the relationship between a company and its customers and stakeholders went through an unprecedented transformation [HTWF12]. For the first time, customers could engage in a near real time manner with companies and brands [SSL13]. The advent of SM radically changed the way customers engage with service providers or product vendors. Any customer could express in an unfiltered way his/her opinion about a brand, a service, a price increase, etc. and the result of it was publicly available in a near real time manner to other customers or customers-to-be. Thus, one could say that the *killer application* of SM in the consumer market has been the customer empowerment. The customer feedback, that used to be trapped in the traditional offline *word-of-mouth* modus operandi, is now available to each and every user willing to know more about the quality of service of any company in the world. SM made these communication barriers fall and changed the customer-company engagement rules for ever, as different types of business are using customer data for better comprehension on customers data [MD12].

As Internet became pervasive with the advent of mobile and wireless technologies –such as Universal Mobile Telecommunications System (UMTS), Long Term Evolution (LTE) and Wireless Fidelity (WiFi)–, posting SM updates or consuming SM content was no longer limited to those sitting in front of a PC with wired access to the World Wide Web. In certain way SM experimented a new (maybe minor) revolution when the access to the Internet escaped the confines of the desktop. Mobile connectivity certainly took SM to a whole new level and brought Twitter's mission statement even one step closer to its realization by making the "instant" component actually feasible. As a consequence of that, the location where the interactions took place increasingly became an integral part of the SM dialogue. New functionality was experimentally launched so that with the user's consent, to each and every SM interaction a place-stamp -in form of a pair of geographical coordinates, the name of a place, etc- could be added. The geo-tagging of the SM interactions started to be supported by the traditional SM platforms and new platforms emerged, where the role of the location surpassed the content itself, such as Foursquare [4], that provides personalised local search experience for its users by taking into account the places a user goes, the things they have told the app that they like, and the other users whose advice they trust. As a result, the proportion of SM interactions that in addition to the known *time-stamp* got a *location-stamp* started to increase drastically, opening the door to a whole new set of insights for a location analysis based on the SM users and the SM interactions tagged in the location [CCL10, CML11]. The accuracy of

[3] https://blog.twitter.com/2012/bolstering-our-infrastructure

[4] https://foursquare.com

the geo-location tags could vary from a few meters in the case of Global Positioning System (GPS) powered pair of latitude-longitude geographical coordinates to the name of a district, a known place or even a city, supporting different kinds of analysis.

There are several ways of attaching a location to a SM interaction, varying in accuracy –level of precision– and granularity –geographical scale–. The most popular SM platforms offered different functionality for both interactions location tagging and querying by location. In Fig 2 we can see for example how two popular SM platforms, like Instagram [5] and Twitter implement the geo-location enrichment at content posting time. Both offer the possibility to attach a high-accuracy location-stamp delivered by the mobile phone's inbuilt GPS in form of a pair of latitude and longitude coordinates or by any location service by looking up the Access Point mac-address into a location database if connected via WiFi [CK02].

The combination time-stamp and place-stamp introduces new analytical possibilities that were not feasible before. The content generated in a location, the community of users authoring this content or just being exposed to this content, the interaction between users related to the location or with other users in a different place, the relationship between users and the engagement with the content, to name a few, might reveal a lot about the character, the behaviour, the preferences and particularities of the people related to this location. The near real-time character of the SM networks adds on top the capability of monitoring for potential changes in metrics enabling therefore a potential fast reaction or even anticipation. In other words, the near-real time character turns SM into the perfect information feed for early warning systems which usually continuously monitor a certain metric and trigger alerts when it becomes higher than a predefined threshold (e.g.: to react to natural catastrophes [SOM10a, YP11, WT02], in the marketing context of financial monitoring [GKR00, CDH12] and in the last years increasingly supporting marketing use cases, as we are going to show in the course of this thesis)

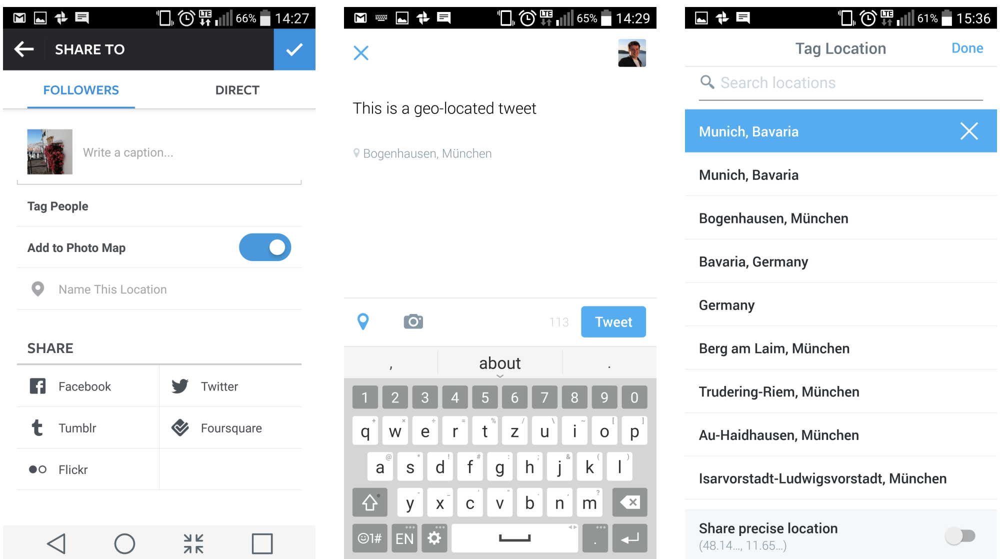

Figure 1: (a) Instagram geo-location feature for uploading pictures (b) Geo-location enable Tweet (c) Named location tag vs. precise location sharing in via GPS coordinates in the Twitter mobile app

The purpose of the work presented in this thesis consists of developing methods to *discover*

[5]`https://instagram.com/`

knowledge based on the SM activity located in an area during a period of time and the information available about the authors of this SM activity (be it derived from the Social Network they are part of or inferred from their behaviour towards other users). The kind of insights produced in this research aims at supporting real use cases that are relevant for the industry, focusing on making insights measurable that have not been quantifiable in an industry-ready way before. Companies can ultimately make use of these insights to steer commercial or rehearsal activities and take action upon.

The output of each method is provided as a set of metrics designed to measure a particular variable over time. Each metric serves to diagnose a problem, to support a decision-making process involving the creation of an action plan and to measure the performance of the actions defined in this plan. Apart from this *traceability over time*, our metrics support the *comparison* of different locations, no matter how different in size or activity levels they are. The activity bias has been carefully removed by design; thus, a place with overly active users can be compared with a place with much lower activity and number of engaging users, as long as the activity does not go down to levels where the volatility renders the metric non-usable. This bias removal does not compromise the *sensibility*, reflecting each particular change upon proper threshold definition. Last but not least, the nature of SM enable the *near-real time character* of the produced insights, which is essential to set up early warning systems, as we said before. Depending on the importance of the real-time character, some methods offer two different working modes: *fast-delivery-low-accuracy* vs. *slow-delivery-high-precision*.

In essence, the four knowledge discovery approaches presented in this dissertation work according to the same pattern: gathering of geo-located SM interactions harvesting, classification, insights enriching inferring knowledge from the social networks the authors of these transactions belong to and a final step of metrics computation.

The performance of each and every knowledge discovery method presented in this thesis has been carefully evaluated in real locations with real SM interactions –in most of the cases tweets–. For that, each set of metrics is defined in conjunction with a specific purpose knowledge discovery system. Each implementation relies on different elements to address the method's particularities, but all of them share a common high-level architecture consisting of a set of common components:

- a *Harvester* polls the SM API for tweets geo-location in the location,
- a *Classifier* flags the tweets according to the particular knowledge discovery method,
- a *User Data Collector* gathers the required information to perform some Social Network Analysis (SNA) tasks
- and finally a module in charge of producing the resulting metrics.

As far of our knowledge, until the writing of this thesis, the methods developed for extracting knowledge about locations exploiting the SM streams didn't comply with the traceability, comparability, bias removal, sensibility, real-time and accuracy vs. time-to-results requirements we introduced in the previous paragraphs. Moreover, to our knowledge, our work is the first attempt to quantify the impact of topics and events on locations with the knowledge extracted from SM and the first one to leverage SM feed to monitor marketing activities for own and competitors campaigns for a defined geographical area.

After this introduction, the rest of this work is organized as follows: Section 2. is intended to provide all the background information relevant for our research. Next, the justification of this memory will be given in Section 3., describing the particular problems we addressed. The overall objectives pursued in this thesis to tackle these problems are described in Section 4.. Section 5. presents a summary of the works that compose this memory. A joint discussion of results is provided in Section 6., showing the connection between each of the objectives and how we have reached each of them. A summary of the conclusions drawn is provided in Section 7.. Finally, in Section 8. we point out several open future lines of work derived from the results achieved.

The second part of the memory consists of four journal publications, organized into two main sections, as listed below:

- Impact quantification methods on locations:
  - A new model to quantify the impact of a topic in a location over time with Social Media
  - Quantifying the emotional impact of events on locations with Social Media
- Geo-localized Campaigning and Quality of Service Monitoring:
  - CARESOME: A System to Enrich Marketing Customers Acquisition and Retention Campaigns Using Social Media Information
  - Leveraging Localized Social Media Insights for Industry Early Warning Systems

Introducción

El mundo digital es hoy en día inpensable sin la exsitencia de Redes Sociales o Social Media (SM). SM se está convertiendo cada vez más en una parte importante de nuestras vidas. La visión estratégica de la comañía detrás de la plataforma Twitter captura en una frase la medida en la que SM ha revolucionado la manera en al que creamos y consumimos informacion: *"Nuestra misión: dar a todo el mundo el poder de crear y compartir ideas e información de manera instantánea, sin barreras"* [6].

SM es un registro digital de fragmentos de la vida diaria, y en su conjunto, se podra considerar como un documento vivo de nuestra cultura; incluso la Librería del Congreso de Estados Unidos está archivando todos los tweets enviados por todos los americanos desde que Twitter empezó a operar. Lee Humphreys, un profesor de comunicación de la Universidad de Cornell en Nueva York, dijo que "estos breves mensajes online –los tweets– pueden revelar cantidad de información acerca de la cultura donde fueron creados" [7].

El uso de SM es un fenómeno en crecimiento [Whi13]. Consumidores de medios se han ido sucesivamente pasando de medios impresos (clásicos) a plataformas digitales. Como resultado inmediato, la comunicación ha dejado de ser unidireccional, con los roles de *autor* y *lector* bien definidos. Con la llegada de la Web 2.0, la definición de *autor* empezó a difuminarse [O'r07]. La blogosfera le dio la posibilidad a los lectores de contribuir a los contenidos publicados por un autor determinado, lo que enriqueció masivamente la información, añadiendo nuevas perspectivas y puntos de vista... Al mismo tiempo, los medios empezaron a democratizarse... todo el mundo era capaz de abrir un blog y la visibilidad de ese blog en los motores de búsqueda era determinada por la gente que lo consideraba relevante (fuera de la búsqueda financiada) [PBMW98].

Las plataformas de SM basada en el concepto de *micro-blogging* llevó esta transformación al siguiente nivel, dado que cada persona podía ser autor y lector al mismo tiempo. El paradigma *publica-primero, cometa-después*, tan popular en la blogosfera, empezó a pasarse de moda. Ahora, todo el mundo podía iniciar una comunicación, enriquecer una hebra ya existente, saltar de una hebra a otra, ignorar, criticar, compartir contenidos como fotos, videos, etc.

La facilidad para publicar, compartir y consumir contenidos estimuló la adopción de estas plataformas de SM como el sitio para hablar a cualquier hora de cualquier tema. El mejor ejemplo es Twitter, que se ha convertido en la plataforma de comunición por excelencia del mundo digital [KLPM10]. En marzo de 2012 la plataforma contabilizó un total de 140 millones de usuarios activos creando una media de 340 millones de tweets al día [Ben12]. La noche del 7 de noviembre, entre las 8:11 y 9:11 pm, cuando el mundo quería compartir los resultados de las elecciones americanas, una media de 9,965 Tweets per Second (TPS) [8] resultaron en la creación de un total de 35 millones de tweets en el espacio de una hora.

Lo que otorga a SM su carácter fundamental, en contraposición a otros canales, es el hecho de que los usuarios expresan libremente –casi sin pensar en ello– lo que pasa por sus mentes literalmente [MZ12] –quizás con la excepción de plataformas con un propósito más definido, como LinkedIn–. Y no importa lo vacío de contenido que parezca un mensaje, cada interacción en SM encapsula una intención subyacente. Una interacción SM puede verse, por tanto, como el resultado de dos decisiones de su autor: el qué comunicar y el cómo transcribirlo en palabras –u otros suportes más ricos, como videos, fotos, etc–. Incluso formas de interacción más específicas de SM, como los

[6] `https://about.twitter.com/company`

[7] `http://www.businessinsider.com/library-of-congress-is-archiving-all-of-americas-tweets-2013-1`

[8] `https://blog.twitter.com/2012/bolstering-our-infrastructure`

"likes" de Facebook o los "retweets" de Twitter están llenas de contenido.

En su esencia, SM empezó como un espacio donde todo el mundo con una cuenta podía interactuar con cualquier otro usuario, compartir un contenido, expresar una opinión, etc. sin estar sujeto a ningún tipo de censura. Como consecuencia de esta democratización en la Web, la relación entre una compañía y sus clientes experimentó una profunda transformación [HTWF12]. Por primera vez, los clientes podían entablar una conversación, por así decirlo, con marcas y compañías [SSL13]. SM cambió de manera definitiva la relación entre clientes de empresas de servicios y estas empresas. Cualquier cliente puede expresar de manera instantánea su opinión sobre una compañía, un servicio, un precio, etc, y esta opinión podía ser leída instantáneamente por otros clientes existentes o clientesd potenciales. Así podríamos decir que la *ventaja competitiva* de SM en el mercado de consumidores consiste en otorgar dicho poder a los consumidores. El feedback de clientes, que en el mundo offline estaba atrapado en el paradigma del márketing de testimonio oral, está ahora disponible para todo el mundo interesado en saber más sobre la calidad de un producto o de un servicio ofrecido por una compañia en cualquier lugar del mundo. SM derribó las barrera de comunicación cambiando la reglas del juego, de tal manera que las compañías actualmente no se pueden permitir no analizar este tipo de información [MD12].

Cuando Internet se hizo pervasiva con la llegada de tecnologías móviles e inalámbircas –como Universal Mobile Telecommunications System (UMTS), Long-Term Evolution (LTE) y Wireless Fidelity (WiFi)–, la publicación de contenidos SM o su consumo dejó de estar limitada al ámbito de los ordenadores de sobremesa conectados por cable a la World Wide Web. En cierta manera, SM experimentó una nueva revolución –tal vez menor–, cuando el acceso a Internet escapó de los confines de las redes tradicionales de primera generación. La conectividad móvil catapultó SM al siguiente nivel e trajo la visión de Twitter más cerca de la realidad, sobre todo en lo que respecta a su carácter instantáneo. Como consecuencia inmediata, el lugar donde las interacciones se crearon se convirtió en una parte integral del diálogo SM. Nuevas funcionalidades se lanzaron de manera experimental, de manera que con el consentimiento del usuario, a cada interacción SM se le podría añadir una marca de lugar –coordenadas geográficas, nombre del lugar, etc–. El geo-tagging de lugares empezó a ser implementado por las plataformas SM más grandes e incluso aparecieron nuevas plataformas donde el papel del lugar se hizo más importante que el mismo contenido de la interacción, como Foursquare [9], que ofrece una experiencia de búsqueda personalizada según en el lugar donde se halle el usuario. Como resultado inmediato de estas funcionalidades, la proporción de interacciones SM con una marca de lugar –a parte de la marca temporal– se incrementó drásticamente, abriendo una puerta a un conjunto de informaciones acerca del lugar basadas en las interacciones SM creadas en ese lugar y en sus autores [CCL10, CML11]. La precisión de la marca de lugar puede variar desde unos poco metros en el caso de Global Positioning System (GPS) a el nombre de un lugar o una ciudad, suportando use cases diferentes. Existen diferentes formas de añadir información acerca del lugar a una interacción, variande en precisión y granularidad geográfica. Las plataformas más populares ofrecen funcionalidad diferente a este respecto. En Fig 2 podemos ver un ejemplo de cómo 2 plataformas muy conocidas, como Instagram [10] y Twitter implementan el enriquecimiento geográfico de maneras diferentes. Ambas ofrecen la posibilidad de añadir coordenadas de alta precisión usando el GPS integrado en forma de par de coordenadas geográficas o mediante un servicio de look up por nombre de lugar, que se determina por el mapeo de la dirección MAC del punto de acceso a una dirección geográfica si se produce una conexión por medio de WiFi [CK02].

La combinación de etiquetas temporales y geográficas abre la puerta a nuevas posibilidades de análisis desconocidas hasta ahora. El contenido generado en un lugar en concreto, a comunidad

[9] https://foursquare.com
[10] https://instagram.com/

de usuarios que crearon o que está expuesta a ese contenido, la relación entre los usuarios y el lugar o de los usuarios y otros lugares, la relación entre los usuarios y su manera de interactuar con el contenido, por nombrar algunos, revelan muchos aspectos acerca del carácter, la conducta, las preferencias y las peculiaridades de la gente relacionada con el lugar en particular. El carácter instantáneo de las redes sociales añade la capacidad de realizar un seguimiento exhaustivo de los cambios potenciales en métricas, posibilitando una reacción rápida o incluso anticipada. En otras palabras, el carácter instantáneo convierte SM en la fuente perfecta de información para implementar sistemas de alerta, que típicamente requieren un seguimiento continuo de una métrica y desatan alertas cuando se supera un tope preestablecido (por ejemplo, para reaccionar ante catástrofes naturales [SOM10a, YP11, WT02], en el seguimiento financiero [GKR00, CDH12] y en los años más recientes, soportando casos de uso en el ámbito de márketing, como vamos a ver en el transcurso de esta tesis)

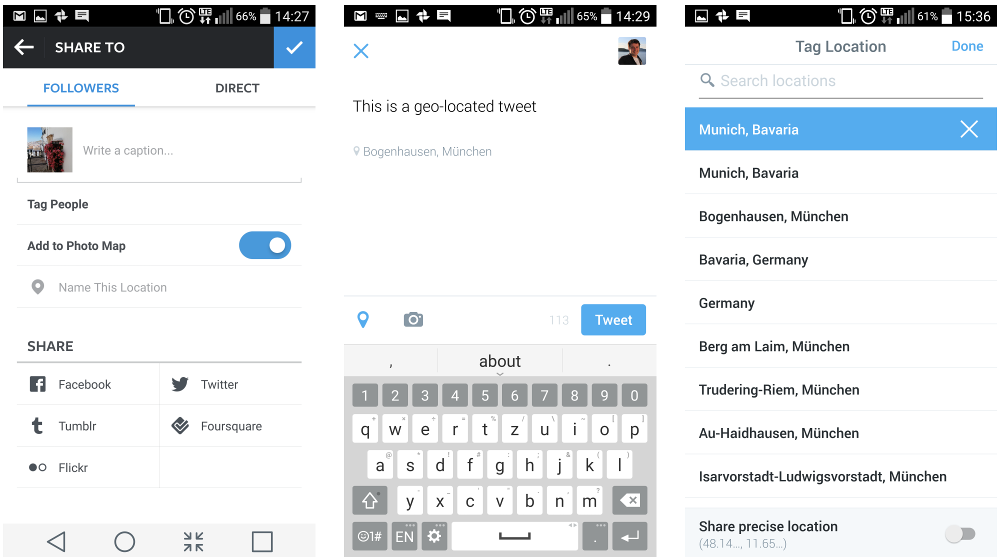

Figure 2: (a) Funcionalidad para la geo-locación en Instagram geo-location (b) Geo-locación en Twitter (c) Geo-locación basada en el nombre de un lugar vs. geo-locación de alta precisión con GPS

El propósito de esta tesis consiste en desarrollar métodos para descubrir conocimiento basados en la actividad SM localizada en un lugar en concreto y explotando la información acerca de los usuarios relacionados con estas interacciones (obtenida analizando la red social de la que forman parte o bien inferida según su reacción al contenido creado por otros usuarios). El tipo de conocimento que producimos en esta tesis intenta suportar casos de uso reales que son relevantes para la industria, haciendo hincapié en hacer este conocimiento medible de una manera que no ha sido factible hasta ahora. Compañías pueden user este conocimiento para gestionar actividades comerciales y de investigación e implementar un plan de actuación.

El resultado de cada método consiste en un conjunto de métricas diseñadas para medir una variable concreta en el tiempo. Cada métrica sirve para diagnosticar un problema, suportar un escenario de toma de decisiones y la creación de un plan con medidas particulares. A parte de esta *trazabilidad en el tiempo*, nuestras métricas suportan la comparación de diferentes lugares, sin importar su tamaño o sus niveles de actividad. El sesgo en la activdad ha sido neutralizado en nuestro diseño; de esta manera, un lugar con usuarios hiperactivos puede compararse con otro donde no exista tanta actividad, siempre y cuando exista un mínimo número de interacciones, si no las métricas se volverían muy volátiles perdiendo valor. El tratamiento del sesgo no compromete su *sensibilidad*, reflejando cada cambio particular basandóse en la definición de niveles de referencia. También añadir que la naturaleza de SM otorga a nuestros métodos la capacidad de medir en una manera *semi instantánea*. Dependiendo en la importancia de esta instantaneidad, nuestros métodos ofrecen la posibilidad de actuar con dos configuraciones diferentes: *velocidad-primero-precisión-despues* vs. *precisión-primero-velocidad-después*.

Todos los métodos para descubrir conocimiento en esta tesis responden al mismo esquema: recolección de las transacciones SM geo-localizadas, clasificación, enrequecimiento de la información disponible infiriendo conocimiento de las comunidades en las redes sociales en las que se hallan los autores de las transacciones y un último paso consistente en la generación de las métricas.

El rendimiento de cada uno de los métodos de descubrimiento de conocimiento presentados en esta tésis se ha evaluado minuciosamente en lugares reales y con interacciones de SM reales –en la mayoría de los casos tweets–. Para esto, cada conjunto de métricas ha sido definido en conjunto con un sistema de descubrimiento de conocimiento. Cada implementación está basada en diferentes elementos para atajar las peculiaridades de cada método, pero todas ellas comparten una arquitectura común que consiste en los siguientes componentes:

- un *Harvester* obtiene los tweets geolocalizados para un logar en concreto mediante la API disponible,
- un *Classifier* marca los tweets dependiendo del método en particular,
- un *User Data Collector* reune la información necesaria de los usuarios relacionados con las interacciones,
- y finalmente un método para crear las métricas de manera agreagada

Hasta donde llega nuestro conocimiento y hasta el momento de escribir esta tesis, los métodos que existían para extraer conocimiento acerca de lugares particulares explotando las interacciones SM no cumplían los requirimientos de trazabilidad, neutralización del sesgo, instantaneidad, precisión vs. velociad, etc, que hemos introducido. Es más, hasta donde sabemos, este trabajo es el primero en abordar la cuantificación del impacto de los tópicos en lugares con el conocimiento extraído de SM y uno de los primeros en consumir la actividad SM para el manejo de campañas a nivel geográfico.

Tras esta introducción, el resto del trabajo se organiza de la siguiente manera: Section 2. está orientada a discutir toda la información de referencia que hemos usado para fundamentar nuestra investigación. Le sigue la justificación de esta memoria en Section 3., describiendo qué casos de uso habilitamos en particular. Los objetivos globales perseguidos en esta tesis se describen en la Section 4.. Section 5. presenta el resumen de las publicaciones que componen esta memoria. Una discusión conjunta de los resultados obtenidos se aborda en la Section 6., mostrando la conexión con los distintos objetivos y cómo se han satisfecho. Siguiendo este punto se presenta un resumen de las conclusiones en Section 7.. Finalmente, en Section 8. hacemos referencia a las líneas de investigación que se derivan de nuestro trabajo.

La segunda parte de esta memoria consiste en las cuatro publicaciones en revistas de investigación de primer cuartal, organizadas en dos secciones:

- Métodos de cuantificación de impacto en lugares:

  - A new model to quantify the impact of a topic in a location over time with Social Media
  - Quantifying the emotional impact of events on locations with Social Media

- Gestión de campañas geolocalizadas y monitoreado de calidad de servicio:

  - CARESOME: A System to Enrich Marketing Customers Acquisition and Retention Campaigns Using Social Media Information
  - Leveraging Localized Social Media Insights for Industry Early Warning Systems

2. Preliminaries

In this section we describe all the background information involved in this thesis. Firstly, Section 2.1 discusses different approaches to knowledge extraction from SM networks. Secondly, Section 2.2 explains the fundamentals of Social Network Analysis and presents the related work we build upon in this thesis. Thirdly, Section 2.3 shows a snapshot on topic diffusion in SM Networks and explains different approaches to impact modelling in SM. Finally, Section 2.4 deals with the different approaches to extract sentiments and emotions from user generated content —especially in SM networks —.

2.1 Knowledge Discovery in Social Media

In this section we are going to describe how SM has been used to extract knowledge for different purposes, indicating first how SM has become so critical for companies nowadays.

Almost every company relies on SM as a communication channel to push company messages and offers, but also increasingly to obtain unfiltered feedback from both existing and prospective customers. Many studies have focused on different aspects of the SM adoption: Kaplan et al. [KH10] highlighted the need for the integration of SM with traditional media to reach customers more efficiently, while defending the advantages of SM to engage with customers in a time-close and high-efficient manner. Mangold et al. [MF09] built upon the idea of considering SM as integral part of the promotion mix, emphasizing the benefit of a less controlled environment to better understand customers.

Several papers focused on researching the role of SM in business and corporations. Jansen et al. in [JZSC09] analyzed the corporate image impact of all interactions related to a brand created over the Twitter channel. In [LBdM13], Li et al. explained the positive impact of the user engagement over the Twitter company channels on the corporate reputation. In [JSFT07] Java et al. demonstrated how similar intentions foster connectivity between users and community building around brands and institutions. Plenty of studies shed light on how companies shall deal with SM related issues like trust and distrust within online communities [KA13] and protection of user's information[MT13].

SM rapidly moved from being *yet another channel* in the communication strategy of a company to be labeled as a *game changer* to engaging with customers: Hennig et al. [HTWF12] explained how microblogging was shaking traditional business models by increasing the role of product quality, especially reducing the time window where product new adopters didn't have any feedback on the product. Culnan et al. [CMZ10] pointed out the need for brands to create communities to exploit the full potential of the virtual customer environments. In [CY14] the link between SM engagement and profitability of online companies was analysed by Chan and his co-authors. In [RBGH13], Rapp and his co-authors analysed the role of SM from the seller, retailer and consumer perspective, demonstrating the value of the SM interactions for better conversion rate.

The effect of the Worth-of-Mouth (WoM) marketing has been extensively researched together in the SM context. Chevalier and his co-authors analysed in [CM06] the effect of book reviews. Villanueva et al. [VYH08] researched the differences in terms of loyalty and equity of customers being acquired through marketing-induced activities vs. WoM gained customers, pointing out performance differences. Bolton established back in 1998 [Bol98] a modeled based on the link between customers retention and customers satisfaction and Rishika et al. [RKJB13] empirically proved the effect of increased SM engagement on the customer visit frequency and customer value.

The need for geo-localized systems to monitor the customer satisfaction at a local scale and to assess the impact of customers' interactions with the brand over SM, emerged [KKT03]. Early warning systems —the equivalent in other domains— have been increasingly adopted in the field of disaster prevention as the sensorial technique allowed for semi-automatic monitoring. There are countless applications for early detection of earthquakes [YP11, SOM10b, GBG11], pandemics [CCMS10, TDRM+06], flood and other natural hazards [Bas06]. In the financial domain fast alerting system have been employed for a wide range of purposes: for example, all variety of economic indicators have been used at a macro level to assess the vulnerability of emerging and existing markets [ZFG11, CDH12, SGLŽ14] and to detect financial crisis in their early stages [BF06, JLCD14], but also at a much more micro level to detect for example critical transactions [SBB+09], etc.

Predicting (i.e., customers) behaviours in SM for management decision making is still challenging tasks [Del03, Liu06, BFK12]. The analysis of SM content and engagement to predict upcoming events has been also intensively researched. In [BAM10] the Bothos, Apostolou and Mentzas explain how agents constantly analysing social media content according to the Belief-Desire-Intentions paradigm can extract enough sentiments and assessments to enable informed decision making in the markets they operate. In [CG12] Colbaugh and Glass employed a stochastic model for dynamic of the interactions based on the underlying network structure to generate useful predictions about the spread of information.

If something makes the knowledge extraction capabilities underlying in SM unique, is the possibility of obtaining it *as it happens*. This *near real time* character itself has been object of research: in [NGKA11] and [PCKJ12], the spreading of bad news takes place really fast over the SM channel, which corroborates their value for the promptly detection of customers' complaints, service outages, etc.. Countless papers built upon the fast news spreading aspect of SM, delivering promising results not only in the industry domain, but also in other areas, such as disaster prevention and crisis coordination: in [SOM10a], 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 in [Cul10] a method to detect epidemic expansion on early stages. In [MMM14], Middleton and his co-authors present a near real time system to map crisis based on several geo-localization techniques of SM information. In the same research line, Yin et al. in [YLC+12] present a system that implements text mining and natural language processing (NLP) techniques to extract situation awareness information from Twitter to support crisis coordination and emergency response. The US Homeland department pioneered the usage of SM to collect real time information about incidents, quantify their extent, monitor their evolution and channel the proper response —programme SMART-C (SM Alerts and Response to Threats to Citizens)—[ASS12].

2.2 Social Network Analysis and Communities

The knowledge discovery methods we apply all along this thesis exploits the static and dynamic properties of the social networks of the users generating or being exposed to content in the particular locations. For that, we apply SNA techniques.

A social network G is basically a mathematical graph [BM76], that consists of nodes or vertices V –or entities in our case– connected by links or edges E, if there is a relation between them; the usual notation is $G = (V, E)$. V might content entities different types of entities, like tags, videos or users, which allows for partitioning V in different subsets and E in different types of edges, depending on the entities they link.

Intuitively, the concept of community relies on grouping nodes so that the density within the group is higher than outside the group [For10]. Long before the existence of social media communities, the community forming patterns have been object of research in Humanities [C+64]. The pioneering analysis of network community structures dated from 1955, where the separation in-group vs. out-group contacts together with other factors defined the groups[WJ55]. Countless sociometric analysis followed to this one, especially in the Sociology domain [Hub65, Bar69, Tic73, RJ71]. In [Gra73], the author pointed out the limitations of the methodologies so far to discover patterns on how small groups aggregate to form large-scale networks and [GN02] exported definitively the graph portioning problem into the mathematical research domain.

At the most abstract level, given a Network $G = (V, E)$, a community can be defined as a sub-graph of the network comprising a set VC subset V of entities or nodes that are associated with a common element of interest [PKVS12]. The transition from the abstract concept of group of nodes to the more tangible concept of community of human entities and their social behavioral attributes arrived as the study of social networks started to be used to understand social interactions [Was94]. In parallel, there have been numerous studies reporting interesting findings about community detection in the Web domain, focusing on hyperlinks as connectivity means among entities [FLG00],[KRRT99]. Several studies proved the correspondence between network communities and building blocks or functional units in more complex real-world networks, for example as pathways in biological networks [GA05], geological units in air transportation systems [GMTA05] or as we mentioned before, websites handling the same topic in the World Wide Web [FLGC02].

Network Communities have been object of research in two dimensions: based on depth analysis or how communities can aggregate into bigger ones establishing a hierarchy, and based on breath analysis [GN02], or identifying which non-overlapping communities or modules can be a meaningful partition for a network. The concept of modularity relies on the homonymous function introduced by Newman et. all [New06]. Community detection algorithms focus on just a part of the network, whereas community structure analysis pursues the holistic view of the network structure[New04].

There have been many methods for communities hierarchy discovery suggested, for example, [SCCH09] relied on histogram analysis of a fitness function applying local optimization for the discovery of the hierarchical structure. In [ABL10] the authors redefine the concept of community as a set of links instead of nodes to address the overlapping issue in hierarchical communities and show the presence of hierarchical structures mapping to geographical division (district, city, region) but also indicate the presence of pervasive overlap. To provide the degree of overlapping between 2 communities, the Jaccard index [Jac01], that proved to be long before the advent of social networks a good overlapping measure in the Biology domain between species [SSP96, LH81] has been broadly adopted as a measure of communities overlapping in the structure analysis of social networks [BGK+05, GKMI+10].

The characterization of communities often relies on the same metrics employed in the SNA, which we are going to refer to in all four publications. The first formal analysis of the role of geographic and psychological proximity in the creation of relationships: *Propinquity* was suggested in the Psychology domain in [FSB50] and then introduced into the online world in [PR05] and connected back to the physical reality in [LLN+95], where the effect of physical distance on the forming of online relationships has been studied. A similar concept introduced in [MSLC01] and adapted to the SM domain in [ABS+12], *Homophily*, quantifies the extent entities tend to form relationships with similar entities vs. dissimilar ones. *Reciprocity*, a property to measure the degree to which two entities reciprocally establish a connection or interact to each other, has been the subject of countless studies in the modern online social networks, for example [MMG+07]. *Triadic closure*, made popular in [Gra73] and thoroughly explored in [FRG10], points out the fact that for three given entities A,B,C, the existence of two strong ties between 2 of them A-B, B-C, implies a weak tie with the third one A-C. The Cluster coefficient [HL71] has been suggested to quantify it. A set of metrics grouped under the concept of Centrality have been developed to assess the importance of an entity within a social network, like betweenness centrality [EB99], Eigenvector centrality [Ruh00], closeness centrality [Sab66], Katz centrality [Kat53], etc. *Density* refers to the proportion of existing direct ties vs. all possible.

Another important concept some of our methods rely upon, is what we call *Tie-Strength*. There have been several studies showing how the Tie-Strength between two SM users plays an important role in the perception of SM interactions. Marsden and his co-authors in [MC84] back in 1984 laid the foundations for measuring the Tie-Strength after Mark Granovetter introduced the concept in 1973 in his paper *"The Strength of Weak Ties"* [Gra73]. In [GK09], a model to predict tie strength by exploiting social media interaction parameters is discussed. The work done by Haythornthwaite [Hay01] confirmed that more strongly tied pairs communicate more frequently, maintain more and different kinds of relations and use more media to communicate. Grabowicz et al. in [GRM+12] analyzed the relationship between SM links and real-world tie-strength and Pan et al. in [PS12] attempted to quantify the role of tie-strength plays in scientific collaboration networks. Shin et al. presented a method to quantify the degree of user sociability in SM relying on the tie-strength[SL12].

2.3 Topic Diffusion and Social Media Impact

Understanding the diffusion patterns of a particular topic, event or news over SM has been key in the development of some of our impact quantifying methods (especially modelling the contribution of a given user based on how active she/he is, the handling of the variation over the time of the topic-related activity and the semantic definition of the topic).

The diffusion of news or topics in the social networks has been subject of intense research especially in the last years [CHLM10, Cen10, SDX13a]. Guille et al. [GH12] defined three dimensions playing a role in the propagation of a topic: social, semantics and temporal to model the probability of dispersion. The social dimension is defined taking into account the users' activity index, the ratio of directed tweets to the user, the mentioning rate and whether the user being mentioned is directly related to the mentioner. On the other hand, the semantics is based on the presence of a keyword in the message being propagated. The temporal dimension is provided as a computation of the user activity in 6 partitions of the day, but probably leaving the door open to finer time granularity.

Rajyalakshmi et al [RBDT12] demonstrated the role of the strong links in the virality of the topics by modelling the diffusion with a stochastic approach, identifying as driving parameters the users activity time and the fading out effect –represented as a weight decay for a topic as time passed by—. In their work, two cases are clearly separated: users creating instances of a global topic or users copying it from their network –local social network effect vs. the overall trending effect—. Romero et al. [RMK11] established a mechanism relying on exposure curves to quantify the impact exposure to other users in making them adopt a new behavior (e.g.: turning them from passive to active contributors or to start using a hash-tag, etc). In addition, there have been several approaches to model the influence of a particular user in his/her own and in the global social media network. Ye and Wu [YW10] defined 3 different metrics to quantify the social influence: followers influence –the higher the number of followers, the higher the influence–, reply influence –the more replies one user receives, the more influential the user is– and re-tweet influence –the more re-tweets, the more influent–. Kwak [KLPM10] suggested also 3 metrics but substituted the reply influence by one inspired by the Google Search PageRank algorithm [PBMW98] to allow the propagation of influence. Depending on the metric applied the ranking of the top users varied. Romero et al [RGAH11] demonstrated that influent users are those whose contributions are not just consumed but also forwarded and therefor overcome the so called passivity and more interestingly, that the popularity of an user and its influence don't quite often correlate. Cha [CHBG10] differentiated 3 kinds of influence for a social media user: due to the size of the user's audience or social network –indegree influence–, due to the generated content with pass-along value –retweet influence, which is also aligned with the passivity– activity work presented by Romero [RMK11] and due to the engagement in others' conversation –mention influence– and all of them are present as component for either Exposure or Engagement when applicable in our approach.

2.4 Polarity and Emotions Extraction from Natural Language

As we mentioned in the introduction, the second pillar together with the exploiting of the insights obtained with SNA, is the analysis of the user generated content available in each and every SM transaction associated to the location under analysis. Apart from the topic, event or brand definition and characterization, we have made use of the two different word mapping techniques: polarity in the context of sentiment analysis and emotional extraction. Basically, both techniques work according to the same schema: each word present in the SM interaction content is looked up in a standardized mapping file and if there is a match, a value or a set of values with a particular meaning are retrieved.

Sentiment Analysis applied to SM has been subject of prolific research. The ground work derives from all the previous studies of terms polarity in the Natural Language Processing domain [WWH05, Loe00]. Pang et al. in [PL08] set the basis of opinion mining based on the analysis of sentiments. Kouloumpis et al. in [KWM11] and Agarwal et al. in [AXV^{+}11] provided an extensive research on sentiment analysis applied to microblogging messages. In this work we rely on SentiWordNet 3.0, implemented by Baccionella et al. in [BES10] on top of the WordNet lexical English database [Mil95].

Emotional models and affective architectures have been intensively researched in the last 15 years in all variety of fields, such as Artificial Intelligence, Human-Computer Interaction, Robotics, Gaming, etc [Hud08]. Yet the first attempts to create a model to compare emotional states were made in the cognitive sciences domain. At an early stage of development, the intensity –or arousal– , the degree of pleasantness –valence– and the amount of influence you feel the environment has upon you –dominance–, were explored independently and represented with different scales. Based on the work initiated in [MR74] and [Meh80] where the *Pleasure-Arousal-Dominance* (PAD) model was formally introduced, Russell suggested in a seminal work the combination of emotional axis to create a circumplex model that enabled the position of emotions on a plane [Rus80]. For the representation of each emotional state, Russell suggested a pair of coordinates on a two dimensional space: on the x-axis the valence and on the y-axis the arousal of the stimulus. Up to 28 emotional states have been multidimensionally scaled in Russell's model, so that intermediate terms are polar opposites (e.g.: excited-depressed, distressed-relaxed, etc). Several new models and refinements on Russell's model followed, each one conceptualizing the dimensions in different ways: tension and energy [Tha89], positive and negative affect [WT85], approach and withdrawal [LBC98], etc.

Bradley and Lang created in 1999 [BL99] a set of normative emotional ratings for 1034 commonly used English words, also known as the set of Affective Norms for English Words or ANEW. Based on the outcome of this research, it was possible for the first time measuring natural language fragments in terms of the PAD model dimensions. This seminal work can be considered the first enabler for the emotional states extraction from user generated content. Fourteen years later, an extended version of ANEW (eANEW) containing more than 13K English lexemes and faceted by gender and education level was developed applying almost the same procedure as in the original piece of work [WKB13]. ANEW became very popular in the research community and was adapted for other languages [VCK+09, RFPC07, SCP+12, WWL11]. In addition to ANEW, further affective dictionaries have been created, for example Affective Wordnet [SV+04], where semantic synsets are assigned one or several affective labels for those concepts representing moods, situations eliciting emotions, or emotional responses.

The recent years have witnessed the creation of countless approaches to extracting emotional states from user generated content. In [RH11] Ramaswamy et al. created an interactive tool to visualize the emotions extracted from a Twitter query over the Russell's 2D plane for the most recent time. This tool also allows for a keyword extraction based on frequency, as well as the visualization of a moods' heatmap over time. In [GAP+10], the authors went even further and mapped the emotions to a 3D virtual human. An additional interesting contribution of this paper is the color interpretation of the emotions mapping different values of arousal and valence to colors. In [SDX13b] the authors analysed the role of the different emotional states in the information diffusion in SM. In [CSS+11] the authors explored the emotions distribution over 2.5 million post in the BBC forum, analysing the correlation between negative emotions and users activity. In [PT13] a predictive model for blog posts ratings providing the estimated level of valence and arousal of a post on a ordinal scale was presented, also taking as a basis the Russell's circumplex model. In [DHK+11] the authors created near real-time, remote-sensing, non-invasive, hedonometer consuming geo-localized tweets from the Twitter Streaming API. The happiness extraction from the micro-posts relies on an own crowd-sourcing effort where over 10000 words were rated for happiness, instead of adopting the traditional ANEW family.

3. Justification

The combination of Social Network Analysis -explaining who is connected to whom, who is exposed to the content generated by whom- and User Generated Content Analysis referred to a particular location allows for the extraction of unprecedented knowledge about this location.

Therefore there are countless uses cases in different areas and disciplines that either benefit from or are for the first time feasible based upon the insights created by the methods we are presenting in this work. In the following subsections we are going to discuss some of these use cases in different domains.

3.1 Marketing Use Cases

As a matter of fact, the way of running marketing is nowadays moving away from a gut-feeling, shot-gun approach to embrace a more scientific, data-driven audience selection [SBI+14].

The customers' targeting and offer selection algorithms need to be nurtured by proper contextual insights. Moreover, it's critical in order to enable a feedback and control loop over time (e.g.: to apply reinforcement learning algorithms [RO15, SNB+13]), having an insights flow continuous and capable of reflecting changes over time.

The methods we propose in this thesis deliver -to our knowledge- unprecedented insights to support many marketing related decision making processes.

One of the most challenging marketing problems is **measuring the outcome of a campaign**

[McD11]. The Impact quantification methods on locations we present in this thesis allows for measuring the total impact and to quantify this impact from the emotional perspective. Thus, a certain campaign on a location can be measured in terms of impact -users engagement, users exposure, recency and frequency metrics- but also we provide the answer to the question: *how people feel about the campaign?*. And unlike the traditional survey-based approaches employed to gain this kind of insights so far, our knowledge extraction systems can continuously run over time, which is essential to define a baseline to measure against.

Another traditional marketing problem revolves around the decision making speed and the time–to–action [RH14]. The opportunity window to react to an event in the market a company is playing is limited. Companies need tools to implement triaging strategies when customers are complaining about their services or to start a campaign to poach customers from competitors and capture market share when some dissatisfaction is detected (e.g.: a Telecommunication operator experimenting over a longer period of time an outage with the corresponding service disruption might be taken advantage of by the competitors to run a campaign promoting the stability and high quality of their network). Our Geo-localized Campaigning and Quality of Service Monitoring methods provide a continuous monitoring at location level and also valuable information to decide on the need and the extent of a reaction.

3.2 Psychology and Sociology Use Cases

Our continuous feed of insights about the people living in a particular location can be equally seen as a gold mine in the Psychology and Sociology domain.

Different communities can be characterized by their reactions to different events and topics which allows for some affinity mining [SCC$^+$13]. Our insights would enable the measuring of the extent local communities care about high-importance topics and happiness drivers, in other words, we could measure *the importance of what people care about* [Fle12] in different places. Moreover, our knowledge extraction methods allow for monitoring for changes and historic comparison (e.g.: people in London care now less or more about air pollution as compared with a year ago?). The same set of metrics can be applied to report changes in taste and affinity (e.g.: quantifying to which extent people in Manchester like more or less the topic *Football* –one of the example we have chosen to discuss the performance of our impact quantification model–; see the first article in the Chapter "Publications"). In the market research field these insights are equally relevant to for example quantify the perception of a brand over time in a particular location [GWDW09] (e.g.: in the middle of the supermarket chains war in UK, the change in the perception of Aldi vs. Tesco over the last 3 years).

The emotional profiling of locations opens the door to new research lines in the Psychology domain. For example, measuring difference between places, understanding what motivated these differences and tracing changes over time. At individual level, it's also very useful to understand how a particular person fits into a community by comparing their emotional profiles and unveiling certain adaptability issues [Nou13].

3.3 Politics Use Cases

Politic Sciences can benefit in from our insights in the same way we explained before for Sociology: getting a deeper understanding about the addressees of a message to tailor it accordingly [HN12, GW13].

But even the particular case of planning a politic campaign on a location could be taken to the next level: monitoring of what has been said about your party and about the competitors party, knowing the impact of the top 10 topics under debate in the particular community (abortion, unemployment, gay marriage, etc), understanding how people in the location emotionally react to these topics, tailoring the message choosing the words in most accord way to the emotional profile of the location, etc... All of them can make a difference in terms of campaign effectiveness and all of them are powered by the methods we are presenting in this work. To which extent it trespasses moral barriers and enters in the realm of the manipulation is up for discussion.

3.4 Use Cases in Other Domains

In general, the outcome of our methods unlock further use cases in all industries where location insights matter. Retail planning, Real State (e.g.: recommending locations based on affinity), Tourism [LLVHB13, HPB13], etc.

Even for Crisis and Emergency Management it's important understanding the emotional profile of the people of the affected location to deliver the best suitable message to avoid overreactions and panic spreading [YP11, JLA11].

4. Objectives

After the study of the current state of all the areas described in the previous sections, it is possible to focus on the actual objectives of this thesis. They will include the research and analysis of the background fields described before, and the development of advanced knowledge extraction methods to turn the geo-localized SM stream into valuable insights.

More specifically, two main objectives motivate the present thesis: to design, create, implement and analyse new methods (1) to quantify the impact of different topics and events and (2) to support the near-real time monitoring of campaigns and the quality of service, in both cases for a particular location and in both cases exploiting the analysis of geo-localized SM interactions. We are going to add a third objective, which is orthogonal to the previous two and applies to all four methods presented in this thesis and which is also a major differentiator of our work –apart from its novelty–: (3) to produce knowledge discovery systems that deliver operationalizable metrics to support industry use cases, as described in the Section 3..

In the upcoming paragraph, we will elaborate the sub-objectives that form each one.

- To quantify the impact of different topics and events for a particular location exploiting the analysis of geo-localized SM interactions
  - **User engagement with a topic**: to measure and quantify the engagement with a topic of all users related to the location and to make the overall topic quantification for the location dependant on that.
  - **User exposure to a topic**: to measure and quantify to which extent all users related to the location are exposed to a particular topic and to make the overall topic quantification for the location dependant on that.
  - **Recurrency and recency aspects**: to incorporate a timely component, so that the impact quantification is always referred to a given time interval within which recency and

frequency metrics can be defined. It also supports the sub-objective "Comparability" we defined for the metrics, enabling non only the comparison of locations, but also the comparison of to different periods of time for the same location or even the comparison of two locations in two different periods of time (e.g.: impact of April's elections in London vs. July's elections in Edinburgh).

  - **Base-lining**: to be able to compute a long-term baseline value or norm for the locations over time, so that the impact of events can be seen as the difference to the baseline.
  - **Quantification at different levels**: to provide in addition to topic quantification insights, the quantified of the impact at emotional level, answering both questions *how important was the event for the people in the location X?* and *how people felt about it?*.

- To support the near-real time monitoring of campaigns and the quality of service for a particular location exploiting the analysis of geo-localized SM interactions.

  - **Market coverage with multiple players**: the possibility to monitor a selection of or all relevant players in the competition arena to identify own and others' weak-spots and act accordingly.
  - **Leverage User internal behaviour for the modelling**: to take into account how people changed their behavior related to a particular event.
  - **Analyse user interactivity with own network**: to take into account how people externalized their behavioral change to their social network.
  - **Sense of severity to assess need for reaction**: to quantify the severity of a particular problem or incident detected in a location based on the analysis of the SM feed. This severity indicator is required to separate might, from could, from must act

- To produce knowledge discovery systems that deliver operationalizable metrics to support industry use cases

  - **Continuity over time**: each metric serves to diagnose a problem, to support a decision-making process involving the creation of an action plan and to measure the performance of the actions defined in this plan.
  - **Comparability of different locations**: no matter how different in size or activity levels they are. The activity bias has been carefully removed by design; thus, a place with overly active users can be compared with a place with much lower activity and number of engaging users, as long as the activity does not go down to levels where the volatility renders the metric non-usable.
  - **Sensibility to changes**: this bias removal does not compromise the *sensibility*, reflecting each particular change upon proper threshold definition.
  - **Near-real time character**: the nature of SM enable the *near-real time character* of the produced insights, which is essential to set up early warning systems, as we said before. Depending on the importance of the real-time character, some methods offer two different working modes: *fast-delivery-low-accuracy* vs. *slow-delivery-high-precision*.

5. Summary

This thesis is composed by four works, organized into two main different parts. Each part is devoted to pursue one of the objectives, and their respective sub-objectives and all of them aim at complying with the third objective, as described above.

- Impact quantification methods on locations:
  - A new model to quantify the impact of a topic in a location over time with Social Media
  - Quantifying the emotional impact of events on locations with Social Media
- Geo-localized Campaigning and Quality of Service Monitoring:
  - CARESOME: A System to Enrich Marketing Customers Acquisition and Retention Campaigns Using Social Media Information
  - Leveraging Localized Social Media Insights for Industry Early Warning Systems

This section presents a summary of the different proposals presented in this dissertation according to the two pursued objectives (Section 5.1 and Section 5.2, respectively). In each section, we will describe the associated publications and their main contents. In addition, we are going to describe a common system footprint to all articles, to explain how the knowledge extraction is implemented.

5.1 Impact quantification methods on locations

This subsection encloses all the works related to the first part of this thesis, devoted to the analysis, creation and implementation of new methods to quantify the impact of different topics and events on a particular location based on the knowledge discovered in the geo-located SM stream. The Subsection 5.1.1 explains how we tackled the problem defining a model based on the standard Recency Frequency and Monetary framework. The Subsection 5.1.2 deals with the impact quantification at emotional level.

5.1.1 Geo-localized impact quantification of a topic based on SM

In this paper we present a new model built upon geo-localized Social Media interactions to quantify the impact of a topic on a particular location and to monitor how it changes over time. As a foundation for our model, we have chosen the well-known Recency Frequency Monetary (RFM) paradigm [Hug05, Kum08, FBPR10]. The Recency component represents in our adaptation last time a given topic was discussed about, Frequency capture how often the topic has popped up in the SM stream and Monetary or Value measures the perceived impact based on 2 components we also introduce in our research: the concepts of **exposure** and **engagement** of a particular user with a topic to model. Concretely in the industry domain, our new social media RFM model could present a good performance in a variety of applications, ranging from event planning and marketing (campaign monitoring, topic affinity advertising, interest targeting) to market research (media monitoring, geo-located panelists, news impact).

The introduction of the exposure and engagement metrics allows for modelling at user level both passive and active topic impact and allows for filtering and segmentation based on different user attributes as all the metrics are defined at user level. As show-cased with the implemented system and the football analysis, our metrics perform well even in hourly chunks; they are consistent over time (delivering similar results in similar situations in different periods) and easy to understand (as they reflect the nature of the social network but at individual level). In a variety of scenarios or extreme cases, the social media RFM model is proven to be robust always delivering meaningful metrics as discussed and demonstrated with the examples we analyzed based on the system that has also been implemented as part of this paper. One of the strengths of the approach we suggest in this work is the fact that the topic impact comparison is supported in heterogeneous scenarios, for example with different topics over different time frames in different locations.

The journal article associated to this part is:

- Bernabé-Moreno, J., A. Tejeda-Lorente, C. Porcel, and E. Herrera-Viedma. "A new model to quantify the impact of a topic in a location over time with Social Media." Expert Systems with Applications 42, no. 7 (2015): 3381-3395.

5.1.2 Geo-localized emotional impact quantification of an event

In addition to the method to quantify the observable impact explored in the previous Subsection, we created a new approach to quantify the **emotional impact** of an event on a physical location based on the analysis of Social Media interactions that have been geo-located in this location.

We explored different modelling approaches for the emotional profiling of locations adopting the well established Pleasantness-Arousal-Dominance paradigm [MR74].

The methods we propose builds upon following components:

- the well-established (P)leasantness or (V)alence-(A)rousal-(D)ominance emotional state model introduced by Russell, to model emotions [MR74].
- an extended version of the Affective Norms for English Words [BL99], to extract emotions from the Social Media user generated content.
- an evolution of the Russell's circumplex model [Rus80] to map the $[v, a]$ scores to one of the set of named emotional states derived from, such as *Impatient, Hopeful, Amorous, etc.*

To quantify the emotional impact, we first introduced the concepts of **emotional baseline** for a location and **emotional footprint** for an event based on the analysis of user generated content posted over SM in the place under analysis.

We define the *emotional baseline* EB of a location L over a given period of time $\triangle t$ as a valence-arousal-dominance distribution resulting from the aggregation of all interactions' emotional ratings authored by the users in the location during the period of time $\triangle t$. We call *emotional footprint* EF of a given event, to the aggregation of the emotional ratings of the interactions related to this event over a period of time $\triangle t$. Both the emotional baseline of a location and the emotional footprint of an event are defined by the multivariate kernel density function applied to the whole set of $[v, a, d]$ scores gathered over the define time period. A detailed explanation and rationale behind the definition can be found in the paper itself (Publications Chapter).

Our method works in the bi-dimensional $[v, a]$ space to enable the mapping to named moods on one hand, and in the $[v, a, d]$ three-dimensional space to consider the effect of the dominance component on the other hand.

The emotional impact is defined as the difference between both concepts and provided a mechanism to measure this impact at a much finer granular level, namely for each particular existing mood.

To evaluate our approach, we implement a system based on Twitter and discuss the results in different scenarios for three known locations in Great Britain: Edinburgh city center, Chelsea in London and the surroundings of Old Trafford in Manchester. Our analysis focused on quantifying the emotional impact of Nelson Mandela's death and Paul Walker's decease at the beginning of Dec. 2013, which we have carried out with different granularity levels –hourly, daily and bi-weekly– showing in a very thorough manner the performance of our method and uncovering the potential to apply it in real-world applications.

The applications of emotional profiling of locations in general and emotional impact measuring in particular are countless. This kind of insights open a new door to marketing activities (e.g.: choosing the right marketing message that fits best the emotional baseline of a location, identifying the best set of promotional activities based on emotional impact, etc.), tailoring of political campaigns (e.g.: selecting the right wording in the massages and measure the outcome) or at a particular level, even finding the right place to live based on the emotional profile of the potential neighbours and their emotional reaction to events. These are just a few examples of the countless applications of the output of this piece of work.

The journal article associated to this part is:

- Bernabé-Moreno, J., A. Tejeda-Lorente, C. Porcel, Fujita.H and E. Herrera-Viedma. "Quantifying the emotional impact of events on locations with Social Media" Knowledge Based System.

5.2 Near-real time monitoring of campaigns and the quality of service on a location

This subsection encloses all the works related to the second part of this thesis, devoted to the analysis, creation and implementation of new methods to to support the near-real time monitoring of campaigns and the quality of service for a particular location exploiting the analysis of geo-localized SM interactions. The Subsection 5.2.1 deals with the creation of a system to leverage discovered SM knowledge to enrich both marketing acquisition and retention campaigns. The Subsection 5.2.2 accent the near-real time character of our methods to produce an early warning system to monitor own and competitors' service quality decays on a particular location.

5.2.1 Enriching marketing customers acquisition and retention campaigns using social media information

In this paper we introduced CARESOME, a Customer Aquisition and REtention SOcial MEdia supporting system that leverages geo-located SM insights to support both customer retention and acquisition activities. CARESOME turns the SM channels into a sensor that companies can use to understanding the impact of the unfiltered feedback given by their customers and prospect customers, but also to uncover competitors' weak spots and engineer acquisition strategies targeting them.

Our system relies on a framework of metrics (see Fig. 3) intended to quantify what we defined as intrinsic and extrinsic impact, where we modelled the contribution of all potential factors playing a role in the impact perception, such as users's engagement with the brand or topic, the underlying communication purpose per interaction and how the authors of these interactions are connected to other SM users.

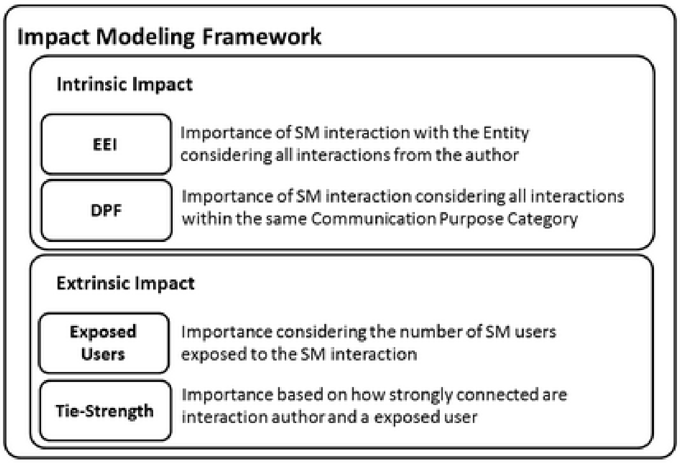

Figure 3: Intrinsic and Extrinsic Metrics metrics overview defined in the CARESOME system

CARESOME is designed to produce actionable insights supporting the customer facing departments of any service company. Thus, in addition to the suggested approach to compute the impact metrics, a speed modus is available, which trades accuracy against time-to-results. To make the generated insights more actionable and enable a prompter decision making, CARESOME also implements a mapping of the results to categories so that the system users do not have to deal with large, hard to compare numbers, but with simple shaded impact categories over time.

The journal article associated to this part is:

- Bernabé-Moreno, J., A. Tejeda-Lorente, C. Porcel, H. Fujita, and E. Herrera-Viedma. "CARESOME: A System to Enrich Marketing Customers Acquisition and Retention Campaigns Using Social Media Information." Knowledge-Based Systems (2015).

5.2.2 Leveraging localized social media insights for industry early warning systems

In this work we present a set of new metrics to measure how a community of customers located in a place is impacted when the quality of the service provided by a company decays and to quantify to which extent the company's image is affected. These metrics are built upon the SM interactions related to the company that are created in that particular location and address several aspects, such as the underlying communication purpose of the interaction, how the authors of these interactions are connected to other SM users and the level of severity, which we computed based on the content polarity.

Our metrics are designed to produce insights to be fed into Early Warning Systems (EWS) for decision making. Therefore, we provide an abstraction layer on top of the metrics mapping their values to levels that can be rapidly actioned by EWS. We also address the cases where the time to results is critical by providing approximations to the single components of our metric and removing therefore the time consuming steps but compromising the precision. The Fig. 4 describes the conceptional knowledge extraction, impact quantification for brand related interactions and severity mapping to the Red-Amber-Green (RAG) scale, which is consumed by the EWS for rapid decision making, as we just mentioned.

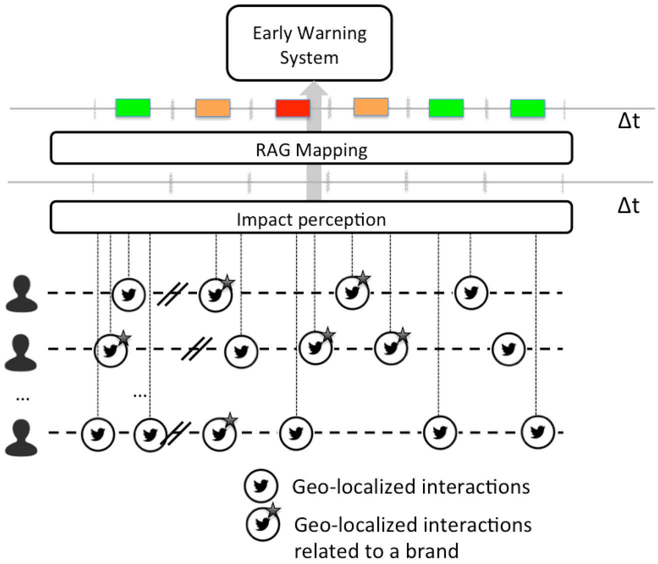

Figure 4: Early Warning Systems feeding from Social Media Interactions

The details about the mapping function and the overall knowledge extraction procedure is given in the corresponding journal paper:

- Bernabé-Moreno, J., A. Tejeda-Lorente, C. Porcel and E. Herrera-Viedma. "Leveraging Localized Social Media Insights for Industry Early Warning Systems" Knowledge-Based Systems (2015).

5.3 The Knowledge Extraction System footprint

The systems we have created to implement the different methods suggested in this thesis present substantial differences. Yet, they all share a common footprint (see Fig. 5). In this Section we wanted to describe this footprint and explain the different design decisions we have made.

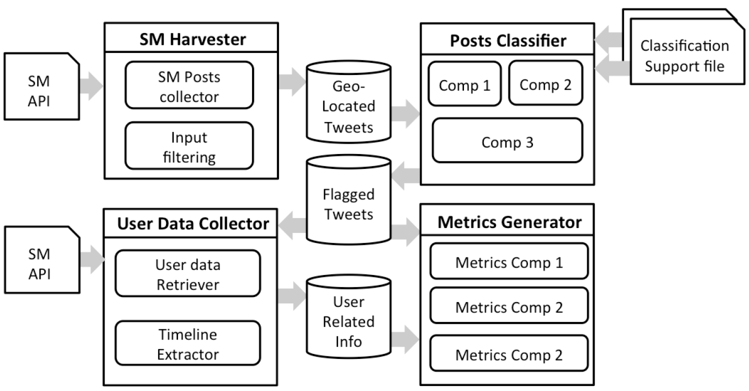

Figure 5: Generic Social Media Knowledge Discovery System Architecture

The Knowledge Discovery methods suggested in this thesis can be applied to each and every social media platform — provided there are means of getting access to the required information—. We've chosen Twitter to implement our system because of following reasons:

- Ease of information extraction: almost no restrictions to get a significant sample of all interactions providing a set of query parameters (using the publicly available Twitter Search API [11]).
- Text-based content dominance: unlike other platforms favouring more rich media content –videos, pictures, etc–.
- High share of geo-located interactions.
- High-engagement general purpose platform.

The system consists of 4 different modules in charge of different labours all along the process.

5.3.1 Social Media harvester

This module performs poll-requests from the Twitter Search API to store the tweets into a local data base for further processing. The tweets are selectively picked for a given area which is configured in the harvester, namely the one we want to extract knowledge from. Additionally, the Twitter API supports the filtering by language (e.g.: only tweets in English), but even it would make the later processing much easier, it might disregard the interactions of all users related with the topic in the target area for being in a foreign language. We opted for a work-around that doesn't filter out the tweets by language upfront, yet doesn't introduce the need for applying NLP techniques in all identified languages.

The harvester also provides the capability of assigning the interactions to already standard geographical output systems (like postal sectors, output areas, census units, etc depending on the country). This is essential to understand the impact of a topic on a location based on other attributes available about the people of this location (e.g. socio-demographic information available at output area level in UK, etc.). The geo-coding module turns the geographical information available in the various fields in the tweet like location, message or geo-coordinates into a pair of latitude-longitude coordinates.

A gazetteer supports both correction of inaccurate information and spelling mistakes in addresses, etc. To implement these functionalities our system relies on existing geo-coding API's provided by the major web mapping providers[12] -which usually are free of use up to a limit of request per day-. The geo-mapper component takes as input the shape files (polygon lines) describing the output geographical units of a system of choice (e.g. postal sectors) and applying the standard point-in-polygon algorithm [SSS74] establishes the mapping of the interaction or tweet to a standard geographical unit. The outcome of the harvester is a collection of full-fledge[13] tweets with a time stamp, a pair of geographical coordinates and potentially the link to a particular standard geography unit.

[11] Available at `https://dev.twitter.com/docs/api/1/get/search?`

[12] Google Geocoding API: `https://developers.google.com/maps/documentation/geocoding/` The location API provided by Bing Maps REST Services: `http://msdn.microsoft.com/en-us/library/ff701715.aspx`

[13] Full-fledge because all the meta information coming from the Twitter API has not been discarded (see `https://dev.twitter.com/docs/platform-objects/tweets`)

5.3.2 Posts Classifier

The mission of this module, in the most generic way, is basically applying some tagging to the harvested SM interactions. Different components perform different tasks, according to the specific requirements of the knowledge discovery method.

Tagging usually works with the so called "Definition files", containing the necessary information to characterize a topic, an event, a brand, etc. These files are typically:

1. *Social Media Entities related to the topic/event/brand*: Set of official accounts, nicknames, hashed tags, etc. users mention in their interactions with the event. For completeness it should include both official accounts and those that are not official but with high levels of activity.

2. *Topic/Event/Brand Named Entities*: set of named entities related to the topic.

3. *Topic/Event/Brand Lexicon File*: containing the set of non-named entities related to the topic.

The tagging step relies on some NLP processing: each geo-located tweet is tokenized applying a sentence tokenizer first and a word tokenizer later (based on [OKA10]) both adapting the Punkt Tokenizer [KS06] to deal with social media texts. The modified tokenizer provides the stop words removal as well. The *event flagger* intents to match each and every reference term listed in the SM and Named Entities files applying a string similarity algorithm [YYZ$^+$03], which delivers a similarity score. The matching procedure implements thresholds —that may differ depending on the source— to support the fact that the social media content is often full of spelling errors [CRA10], which is likely to happen even more frequently when it comes to named entities of foreign people.

In our CARESOME system, there are also definition files for each Communication Purpose Category, whose tagging takes also place here. For our method to quantify the emotional impact, the Mood mapping and the VAD rating take place in the Posts Classifier too.

5.3.3 User Data Collector

This component is in charge of polling the Social Network of all the authors of the tweets identified by the first pass of the classifier. Additionally, the harvester comes again into picture to retrieve the tweets of each and every user in each author's social network to enable the exposure analysis. Once all tweets of all users in the social network for the time-period being analyzed have been stored, the classifier acts again to flag those belonging to the topic/event/brand.

The segmentation of the users' base is also supported and carried out by this component. Examples of segments could be: users socially well connected, users settle in the target area, users who are engaging with social networks for more than X years, users with higher activity index, etc.

The User Data Collector component supports the metrics comparison in an area over time for different segments, which certainly helps understanding the impact drivers and defining predictive models based on the segment attributes.

5.3.4 Metrics Generator

As a final step, the *Metrics Generator* pulls all the metrics together to generates the different specific purpose output metrics. For that, the set of equations specified in each knowledge discovery method is applied. Usually, the metrics are first compute at user level and then aggregated up using a parametrizable function.

If the system offers a *fast-delivery-low-accuracy* mode, a component in this module takes care of creating this set of metrics.

6. Discussion of results

The following subsections summarize and discuss the results obtained in each specific stage of the thesis. We are going to revisit the objectives we formulated in the Section 4. and discuss whether we managed to meet all of them in the course of this dissertation.

6.1 Impact quantification methods on locations

In this subsection we are going to share the results obtained in the experiments we performed to evaluate our knowledge discovery approaches for impact quantification and how these results underpin the objectives of this thesis.

6.1.1 Geo-localized impact quantification of a topic based on SM

In order to evaluate how the suggested RFM model behaves in different scenarios to prove both sensibility and usage of the set of defined metrics, we performed a set of experiments.

Our validation proposal consisted of two real world topics to ensure the coverage of all the variety of events a topic might manifest.Our first topic revolves around two events that shook the hearts of multitudes within the space of one week: the deaths of the famous American actor Paul Walker 5 on November the 30th 2013 and the decease of the charismatic Peace Nobel Price winner, South Africa first black president and anti-apartheid icon Nelson Mandela 6, just 6 days later. These one-off events are going to help us demonstrate the role of recency and frequency taking different time analysis windows (centered on the day of the death, the week after, the week before, etc). We compared the impact of both deaths in the considered locations, to demonstrate the comparability of our impact quantification metrics.

As second topic we chose Football, much wider in scope but highly suitable to prove the performance of our metrics due to the following reasons: *recurrence* –there are regular matches coming every week–, *variety of scenarios* –team playing at home, as visitor, national championship, Champions League, etc.–, *fine granularity in time* –which allows for engagement and exposure calculations for hourly intervals–, *popularity* –with a lot of social media content generated about the topic and therefore, less volatile– and *easy to model* –with a rather large volume of SM and Named Entities and comparatively small lexicon, which natively reduces the need for disambiguation and therefore, the number of false positives–.

We set up 5 harvesters: 2 of them to monitor the activity on two well-known football stadiums: Stamford Bridge (Chelsea FC) and Old Trafford (Manchester United FC), 2 additional ones centered on both stadiums but with a much larger radius (5km) covering an important part of London and

Manchester and a last engine also with a radius of 5 km covering the city of Edinburgh (a place a priori not so much related with the topic football). Although our harvesters have been running for longer than 3 months, we are going to focus our analysis on the first two weeks of December 2013, where the vast majority of scenarios manifest. The harvesters gathered 1088627 tweets during these 2 weeks in the mentioned locations.

We performed a hour-by-hour comparison of the impact value for the day right after the deaths happened, where the sensibility of our metrics became evident. The comparability was also proven between harvesters but also for different events in different time windows for the same location. For the aggregated view over chosen weeks, we provided each and every RFM component and for the value, we separated the passive exposure from the active engagement to make clear to which extent each component contributed to the resulting impact metric.

6.1.2 Geo-localized emotional impact quantification of an event

To show the performance of out method for quantifying the emotional impact we have chosen 2 tragic events and analysed the results we in three different real-world locations. The event we have chosen are again Mandela's and Walker's deaths. The locations Edinburgh city center, Chelsea in London and the surroundings of Old Trafford in Manchester (each harvester configured with a radius of 5 km). We also focused our analysis on the first two weeks of December 2013 and worked with the 1088627 tweets gathered during these 2 weeks in the mentioned locations.

This choice has been done on purpose, so that the reader has the possibility of understanding both impact quantification methods applied to the same scenarios. Our intention can has also been to demonstrate how complementary they are and the richness of insights that both together can produce for an event and a location.

As explained all along the paper, the pre-requisite for the emotional impact quantification is the emotional base-lining of the locations and the creation of the event emotional footprint. We have obtained them in two time-granularity levels: hourly and daily. Providing a hourly view over time helps us understanding the carousel of emotions that such a tragic event like the death of these two beloved personalities can trigger. We observed already noticeable differences in all three locations.

When we computed the emotional footprints of both events, we saw that the emotions were very changing, that's why an impact quantification makes more sense at daily level; having more interactions related to the event (1 day vs. just 1 hour) makes the analysis results less volatile on one hand and changing emotions get to equalize along the day on the other hand. Nonetheless, the hourly change of the emotional footprint in both cases is of great interest for appreciating the variety of emotions that a tragic incident can release.

Weighting the difference between event's emotional footprint and location's emotional baseline by the number of interactions associated to the event (as suggested in our framework), reports the final emotional impact quantification over time for both events and all three locations: we could say that both deaths have similar impact in Chelsea, but while Edinburgh has been definitively more impacted by Mandela's, Walker's death left a deeper mark in Manchester. The results obtained for Edinburgh show us the role of the event's share in the impact metric: while Walker's death presents a more diverging emotional footprint from the emotional baseline of the Scottish capital, the share is much lower than Mandela's and therefore the overall impact.

Our method also quantifies the particular changes at named mood level. In our experiment, we can say that Mandela's death massively *impressed* people in all three locations. *Expectancy* was also observable in Manchester and Edinburgh, while Chelsea reacted more *contemplatively.*

Remarkable uplift of *apathetic feelings* and people *taken aback* in all locations. Walker's decease released a generalized *discontent* in the English cities. *Apathy* is also noticeable in a general note as well as *expectancy*. Manchester and Edinburgh showed an increase of *distrust* and *discomfort*, while in Chelsea.

6.2 Near-real time monitoring of campaigns and the quality of service on a location

In this subsection we are going to share the results obtained in the experiments we performed to evaluate our knowledge discovery approaches for Near-real time campaign monitoring and quality of service monitoring, and how these results underpin the objectives of this thesis.

6.2.1 Enriching marketing customers acquisition and retention campaigns using social media information

To analyze the performance of our CARESOME system in action we chose 2 well-known locations with a high volume of visitors and where people are likely to have time and therefore prone to create Social Media interactions: the biggest two airports in the city of London, namely Gatwick and Heathrow.

We set 2 harvesters centered in the middle point of both airports with a radius big enough (5km) and we let them run between the 24th of November 2013 and 23rd of January 2014, collecting 852319 SM interactions in total. During this period of time there were severe weather conditions, spreading the chaos all over the country with strong winds and flooding episodes, which impacted the quality of all transportation services in UK. Thousands of passengers were affected and the Social Media platforms filled with users? statements on how well the different carriers handled the incident.

For our show case we took as entities a subset of the airlines operating in these airports. As explained in the paper, each iteration has been assigned to one of more Communication Purpose Categories: Complaints and Criticism (c), Praise and Positive Feedback (p), Information Request and Customer Care (ir) and a fourth one for the rest called Neutral (n).

The adverse weather conditions on the 24th and the 25th of December left thousands of passengers stranded in the Gatwick Airport due to power problems. In this emergency situation, a blame game between Gatwick airport and the airline Easyjet started 4. Especially on the 24th we observe a peak over all metrics, motivated by the increase of SM interactions (297 different users) criticizing the way Easyjet handled the emergency situation. These results produced by CARESOME would have given Easyjet enough quantified evidence to trigger some sort of reaction and the corresponding communication back to the SM channel to palliate the incident effect. After the potential airline reaction, CARSOME can then measure the SM community response. In general, CARESOME's role is providing enough insights for a company to steer the SM dialogue in all fronts.

At the same time, an Easyjet's competitor, Ryanair, also affected by the weather conditions, decided to adjust their flights schedule, so that everybody could get home on the Christmas Eve. For Ryanair, the category *"Praise and Positive Feedback"* showed a tremendous peak in Gatwick, showing that such a decision had a huge impact on the SM channels. CARESOME was able to identify a weak spot for Easyjet in Gatwick and a customers' satisfaction peak for Ryanair at the same time in the same place. Our system in Ryanair's hands or Easyjet's hands would have served for developing attacking or defending strategies even further and to near real-time measure their effects.

During the monitoring period we also showed how CARESOME was able to detect minor things (sensibility). A British Airways captain successfully landing a plane after complicated maneuvering with adverse weather conditions might be seen as part of his job, but might also trigger a set wave of SM interactions praising the action. Something what passed unperceived for British Airways, might have well used to create a campaign to reinforce the idea of security in extreme conditions.

With this test we proved how CARESOME's metrics are fully operationalizable and deliver precious insights for uncovering competitors' weak spots to build acquisition campaigns upon, or to identify own deficiencies, which could lead to retention campaigns. In both cases, CARESOME would be the tool to measure their performance.

6.2.2 Leveraging localized social media insights for industry early warning systems

We have chosen a real world scenario to demonstrate how our system and the impact metrics can leverage SM insights to feed EWS. Our experiment is based on the same set of Tweets collected by our harvesters Heathrow and Gatwick airports, during the same period of time

Instead of analysing the airline industry, our Early Warning System focused on the railway transportation sector because a) the amount of people using trains on a regular basis is significantly large and b) the customer satisfaction is usually low, which pushes 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.

To model the severity in accordance with the selected Communication Purpose, we just considered the negative polarity. The values have been properly mapped to severity levels and factors, as explained in the paper.

We computed the Impact metric for all 6 entities over 2 months in both airports. The highest value originated on Dec 24 for the Entity *National Railway*, reflects the train service disruption when the storm *Emily* was striking the country services suspended due to strong storm on December 24th. The second highest registered also for National Railway on Jan 17 was due again to weather causing flooding. 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.

We demonstrated that our metric could also quantify the impact of small decays in the quality of service. We compared the total delay in minutes accumulated day by day by the First Capital Connect train lines to or over Gatwick airport with our impact metric. The days with high delay values are usually reflected as peaks in the Impact curve for FCC, 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 EWS 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.

As the timely aspect is so critical for Early Warning Systems, we wanted to prove that our metrics are designed to work in finer levels of granularity, whenever the number of SM interactions remains at a significant level. For the National Railway we have created an hourly heat-map (available in the publication), showing the value of hour impact assigned to deciles that we shaded according to the intensity by hour. This hourly metric can be taken to decide on the cases where the EWS shall generate an alert to trigger a triaging plan.

7. Concluding Remarks

In this thesis, we have addressed several problems pursuing a common objective: developing methods to *discover knowledge* based on the SM activity located in an area during a period of time and the information available about the authors of this SM activity (be it derived from the Social Network they are part of or inferred from their behaviour towards other users).

More specifically, three main objectives have been pursued in the present thesis:

- The design, analysis and implementation of new methods to quantify the impact of different topics and events for a particular location over time exploiting the analysis of geo-localized SM interactions.
- The creation and implementation of new approaches to support the near-real time monitoring of campaigns and the quality of service for a location, based on the insights inferred from the analysis of SM interactions attached to this particular location.
- The design and implementation of knowledge discovery systems that deliver operationalizable metrics to support industry use cases based on our methods.

To satisfy our objectives, we followed two main research lines in our dissertation: Impact quantification methods for locations and Near-real time monitoring of campaigns and the quality of service on locations.

In the first research line we presented a model based on the industry standard Recency Frequency Monetary (RFM) paradigm to quantify the impact of a topic based on the engagement and exposure of the users in the location with/to the topic over time and modelling the time component via Recency and Frequency metrics. We built a system and ran different experiments with three real topics on popular British locations to prove the performance of our results as well as the advantages of relying on RFM to model the time component.

In a second paper we tackle the impact modelling from a different perspective. Relying on well-known cognitive sciences procedures for modelling emotions (Russel's Circumplex emotional plane) and to extract emotions from natural language texts (ANEW), we proposed a method to quantify the emotional impact of a particular event on a location. To compute the impact, we first create a long term emotional baseline of the location by processing each and every geo-localized SM transaction and then we calculate the emotional footprint of the event, being the impact the difference of both. We also provide the impact at named mood level (e.g.: event XYZ impacted the mood "anxiety" by 2.8% in the location A). The results obtained for the emotional impact monitoring of two events in three different British locations proved not only the richness of our modelling approach but also the potential to operationalize our insights for long term monitoring.

The second part of this dissertation has been devoted to campaign enriching and service quality monitoring on a location. We presented our CARESOME system which relies on a set of intrinsic and extrinsic metrics, where we modelled the contribution of all potential factors playing a role in the impact perception, such as users's engagement with the brand or topic, the underlying communication purpose per interaction and how the authors of these interactions are connected to other SM users. The metrics are first defined at user level and then aggregated up for decision making and offers a low-accuracy-fast-insights mode, which trades accuracy against time-to-results to allow for quicker reactions. To prove its performance we implemented the system and modelled the airlines industry in two major airports in London and discussed the impact of different events on our metrics over the course of a 2 months. In our test we for example showed how our system identified the need for a customers retention/defend reaction for an airline and the opportunity for attacking/acquiring new customers for the competing one on a location when the weather conditions caused a service disruption.

In addition we presented a system intended to serve as continuous feed for industry Early Warning Systems. The system relies on a framework of metrics designed to measure how a community of customers located in a place is impacted when the quality of the service provided by a company decays and to quantify to which extent the company's image is affected. These metrics are built upon the SM interactions related to the company that are created in that particular location and address several aspects, such as the underlying communication purpose of the interaction, how the authors of these interactions are connected to other SM users and the level of severity, which we computed based on the content polarity. The suitability and performance of our method has been discussed based on the monitoring of service quality for the major railway companies in the London area.

In this thesis we both proved that our work met the formulated objectives and help the reader understand the value of our contribution.

Conclusiones

En esta tesis hemos abordado diferentes problemas persiguiendo un objetivo común: el desarrollo de métodos para descubrir conocimiento basados en la actividad geolocalizada en un área observada en las redes sociales por un lado y en el conocimiento inferido de las comunidades sociales a las que pertenecen los usuarios relacionados con esta actividad por otro (bien fijándonos en su conducta, bien explorando la estructura de estas comunidades).

De manera más específica, hemos perseguido los siguientes objetivos:

- El diseño, análisis y la implementación de métodos para cuantificar el impacto de diferentes tópicos o eventos para un lugar particular en el tiempo, explotando el análisis de interacciones sociales geolocalizadas.

- La creación e implementación de métodos para habilitar el monitoreo semi instantáneo de campagnas y la calidad de servicio para un lugar en concreto, basado en la información derivada del análisis de interacciones sociales relacionadas con este lugar en particular.

- El diseño y la implementación de sistemas para el descubrimiento de conocimiento capaces de proporcionar métricas operacionalizables para el suporte de casos de uso relevantes a nivel industrial.

Para cumplir nuestros objetivos, nuestra tesis sigue dos líneas de investigación: métodos de cuantificación de impacto en lugares particulares y monitoreo semi instantáneo de campañas y de la calidad de servicio en lugares particulares.

En la primera línea de investigación presentamos un modelo basado en el conocido estándard Recencia, Frecuencia y valor Monetario (RFM) para cuantificar el impact de un tópico basado en la interactuación y exposición de los usuarios localizados en el lugar con el/al tópico y modelando la componente temporal usando Recencia y Frecuencia. También construimos un sistema y ejecutamos diferentes experimentos con 3 tópicos diferentes en conocidas ciudades británicas para probar el rendimiento de nuestros resultados y las ventajas de usar el paradigma RFM.

En una segunda publicación abordamos el modelado de impacto emocional. Basándonos en procedimientos conocidos en la Ciencias Cognitivas de modelado de emociones (Plano emocional circumplejo de Russel) y procedimientos de extacción de emociones en textos del lenguaje natural (ANEW), proponemos un método para cuantificar el impacto emocional de un evento particular en un lugar dado. Para calcular el impacto, primero creamos la referencia emocional para el lugar precesando todas las interacciones SM disponibles para luego calcular la huella emocional del evento, siendo el impacto la diferencia de ambos. Nuestro método también proporciona el impacto a nivel de estado de ánimo particular (por ejemplo: el evento XYZ impactó el estado de ánimo "ansiedad" en un 2.8% in el lugar A). Los resultados obtenidos del monitoreo del impacto emocional de dos eventos diferentes en conocidas metrópolis británicas demostraron no sólo la riqueza de nuestro modelo, sino también el potencial existente para operacionalizar nuestras métricas para períodos más largos de tiempo.

La segunda parte de esta disertación se dedica al enriquecimiento de campañas y al monitoreo de la calidad de servicio en un lugar particular. Presentamos el sistema CARESOME que se basa en un conjunto de métricas extrínsecas e intrínsecas, donde modelamos la contribución de todos los factores que potencialmente juegan un papel en la percepción de impacto, como la interacción de los usuarios de las interacciones con las marcas o los tópicos, el propósito inherente de comunicación y cómo los autores de las interacciones están conectados a otros usuarios en las redes sociales. Nuestras métricas se definen primero a nivel usuario para luego aggregarse y obtener un resultado para el lugar que se está analizando, ofreciendo un *velocidad-primero-precisión-después* modo, que renuncia a alta precisión en aras de ofrecer resultados semi instantáneos. Para demostrar el rendimiento de nuestra propuesta, construimos un sistema y modelamos el mercado de aerolíneas en los dos aeropuertos más importantes de Londres, y discutimos los resultados obtenidos en un período de dos meses. En nuestro test por ejemplo identificamos la necesidad patente para una aerolínea de correr una campagna de fidelización de clientes y para otra competidora, la oportunidad de hacer una camapaña de ataque, cuando la región fue afectada por severas condiciones meteorológicas causando una interrupción en el servicio.

Adicionalmente, presentamos un método para alimentar a un sistema de alertas. Nuestro sistema se basa en un conjunto de métricas orientadas a medir cómo una comunidad de usuarios localizados en un lugar es impactada cuando la calidad de un servicio ofrecido por una compañia decae y que también cuantifica hasta qué punto la imagen de esta compagnia queda afectada. Estas métricas se construyen en base a las interacciones geolocalizadas relacionadas con la compañia en cuestión y cubren diferentes aspectos, como el propósito de comunicación subyacente, cómo los diferentes usuarios están conectados en la red social y el nivel de severidad, que se computa usando la polaridad del contenido de las interacciones. La practicabilidad y el rendimiendo del método han sido ampliamente discutidos basándonos en el monitoreo de las grandes compañias de ferrocarril inglesas en el área de Londres.

En esta tesis hemos probado que nuestro trabajo cumple los objetivos establecidos por una parte y hemos ayudado al lector a entender el valor de nuestra contribución por otra parte.

8. Future Work

The results achieved in this PhD thesis may open new future trends in different challenging problems. In what follows, we present some research lines that can be addressed starting from the current studies.

Extrapolation to the entire population / Online Bias removal All four methods produce insights based on the geo-located interactions of SM users in a particular location. To which extent our insights can be held representative for the entire population of the location, including those who are not SM active and those who are not even digital affine is a question that needs addressing. A new research line could tackle exactly that: quantifying how representative our insights are, quantifying the bias introduced by just measuring SM and suggesting approaches to remove this bias.

Improving geo-tagging capabilities The systems to implement the methods we proposed in this thesis, rely on the geo-location capabilities of the Twitter Search API to periodically retrieve all the interactions of any kind created over SM channels in the specified area.

A limitation is that some transactions created by users in the area are not geo-localized and can't therefore be retrieved by a geo-query. To overcome this problem, a potential solution would be implementing a *user-place stickiness factor*, which computes based on the user's history of interactions, the likelihood of a particular interaction to be located in the area under analysis. Implementing such an approach would improve the data gathering recall.

Differential Influence of SM Users Putting Tie-Strength aside, the importance of the user within the social network of the follower could help improving the performance of our proposals. The impact caused by the interaction of certain user on another one depends to certain extent on how important the first one is within the SM network of the later. Ranking users within a SM network requires complex modelling which we have not addressed in any of our works.

Another possible improvement would be introducing a reputation index for SM authors, which can be taken into account in the impact computation on the SM network. Another improvement could be achieved by modelling the quality of the interaction, approach which has delivered good results in the recommender system domain [SGHVO+11, PAMHV12, TLPP+14].

Adding predictive capabilities Our approaches deliver insights that quantify what happened in the past or what just happened (leveraging the near real time capabilities when running on *fast-delivery-low-accuracy* mode). A very interesting extension with countless application could be made by incorporating predictive and forecasting capabilities.

Predictive Analytics in SM has been subject of intensive research [AH10, ZFG11, GH12] and some of the explored methods could help us predicting the impact of an event or anticipating a warning situation for a location, just by applying these models to the SM stream.

Trend analysis and trends changes detection So far the examples we have used to discuss the performance of our knowledge discovery methods have been very short in time. The quantification metrics they produce have the potential of detecting changes in taste, feelings, interests on a location over a much longer period of time.

This potential can open a new research line focused on applying trend analysis algorithms to our methods and even more interestingly, co-joint analysis with other temporal series to understand what is causing this potential changes in trends (e.g.: weather, macro and micro economic indicators, Google Trend [14]search volume series, etc).

[14]https://www.google.de/trends/explore

Chapter II

Publications: Published and Submitted Papers

1. Impact quantification methods on locations

The journal papers associated to this part are:

1.1 A new model to quantify the impact of a topic in a location over time with Social Media

- Bernabé-Moreno, J., A. Tejeda-Lorente, C. Porcel, and E. Herrera-Viedma. "A new model to quantify the impact of a topic in a location over time with Social Media." Expert Systems with Applications 42, no. 7 (2015): 3381-3395.
  - Status: **Published**.
  - Impact Factor (JCR 2013): 1.487
  - Subject Category: Computer Science, Artificial Intelligence. Ranking 29 / 688 (**Q1**).
  - Subject Category: Computer Science, Computer Science Applications. Ranking 59 / 1129 (**Q1**).
  - Subject Category: Computer Science, Engineering. Ranking 21 / 527 (**Q1**).

A new model to quantify the impact of a topic in a location over time with Social Media

J. Bernabé-Moreno[a,*], A. Tejeda-Lorente[a,*], C. Porcel[b,*], E. Herrera-Viedma[a,*]

[a] *University of Granada, Department of Computer Science and Artificial Intelligence, Granada, Spain.*
[b] *University of Jaén, Department of Computer Science, Jaén, Spain.*
[c] *Department of Librarianship and Information Science, Granada, Spain.*

Abstract

Social media can be used as a thermometer to measure how society perceives different news and topics. With the advent of mobile devices, users can interact with social media platforms anytime/anywhere, increasing the proportion of geo-located social media interactions and opening new doors to localized insights. This article suggests a new method built upon the industry standard Recency, Frequency and Monetary model to quantify the impact of a topic on a defined geographical location during a given period of time. We model each component with a set of metrics analyzing how users in the location actively engage with the topic and how they are exposed to the interactions in their social media network related to the topic. Our method implements a full fledged information extraction system consuming geo-localized social media interactions and generating on a regular basis the impact quantification metrics. To validate our approach, we analyze its performance in two real-world cases using geo-located tweets.

Keywords: Social Media Impact, Topic Engagement, Topic Exposure, Social Media Sensor, geo-located Social Media, Social Network Analysis, RFM, Information Extraction

1. Introduction

The usage of social media (SM) is an ever-growing phenomenon (White, 2013). Media consumers are increasingly shifting from classic (printed) media to digital platforms. As a result the communication stops being one-way with clearly defined *author* / *reader* roles. With the advent of the web 2.0, the definition of *author* started to blur. The blogosphere empowered readers to make their own contributions to the content published by a given author, which radically increased the information richness, adding further perspectives and points of view. Simultaneously, media started to be democratized, as anybody could start a blog and the visibility of the blog in the search engines was determined a priori by the number of people that considered the blog to be relevant outside the realm of the paid search (Page et al., 1998).

*Corresponding Author

Email addresses: juan.bernabe-moreno@webentity.de (J. Bernabé-Moreno), atejeda@decsai.ugr.es (A. Tejeda-Lorente), cporcel@ujaen.es (C. Porcel), viedma@decsai.ugr.es (E. Herrera-Viedma)

The SM platforms based on the concept of *micro-blogging* took it to the next level, as everybody could be an author and a reader anytime. The *push-first, comment-later* paradigm so popular in the blogosphere started to look old-fashioned. Rather, anybody was empowered to initiate a communication, enrich an existing thread, jump from a thread to another one, ignore, criticize, share richer content like pictures, videos, etc. The ease of publishing, sharing and consuming content boosted the adoption of these social media platforms as the place to talk anytime about everything with everybody. The best example is Twitter, which has become a communication platform for almost all the digital world (Kwak et al., 2010). By March 2012 the platform counted 140 million active users creating an average of 340 million tweets a day (Bennett, 2012). The night of November 7th, during 8:11 and 9:11 pm when the world wanted to share the results of the US elections, an average of 9,965 Tweets per Second (TPS) [1] resulted in the creation of more than 35 million tweets within one hour.

With the advent of wireless internet technologies based on WiFi hot-spots and mobile communication networks, the social media content creation became more pervasive. The access to the digital media was no longer exclusive to desktops; the rise of the smart-phones and mobile data packages enabled the always-on era and opened the door to a new set of insights based on the location where the user interacted with the social network. As the proportion of geo-located SM interactions increased, the *geo-fencing* or delimitation of the location boundaries where the SM dialogue took place became more accurate. These new capabilities led to more meaningful and representative analytic results, to the point that the SM activity could be taken as a good indicator of what is happening anytime anywhere.

The influence and impact in the SM channel has been matter of research almost from the advent of the modern SM networks and platforms. Yet, the research community mainly focused on understanding and modelling the impact of a particular user or a particular group of users on their own and foreign social networks. Our intent here is to prove that the impact of a given topic can be measured, quantified and monitored over time. Obviously, this topic centric geo-located impact measuring would open a new window of possibilities in different domains, such as understanding the performance of marketing campaigns on a given area, or understanding the affinity of local communities to certain marketing offers. Likewise, the generated insights can be used in the area of recommender systems and applied in different scopes, especially in e-commerce and digital media (Porcel et al., 2012; Tejeda-Lorente et al., 2014). Unlike the metrics typically used to assess the SM influence of a particular user, which mainly rely on well-defined entities and parameters present standard-wise in social network platforms (like *User*, *Friend*, *Follower*, etc), there's no entity to represent a *topic*. Thus, modelling techniques need to be applied, which introduces a new level of complexity entering in the realm of semantic web (Berners-Lee et al., 2001) and Natural Language Programming (NLP) (Manaris, 1998). Although our aim was not solving any NLP problem, we implemented a system to extract the required information from the social media networks and to apply the quantification methodology for a topic which relies on a whole set of NLP components.

The purpose of this paper is to define a new method to quantify the impact of a topic during a period of time on a

[1] `https://blog.twitter.com/2012/bolstering-our-infrastructure`

given place based on how the users located in this place are exposed to the topic over their social network and how they actively interact or engage with the topic themselves. In other words, we want to turn the social media platform Twitter into a topic impact thermometer. Our method relies on the well-established industry-standard Recency, Frequency and Monetary schema (RFM) (Bult and Wansbeek, 1995). RFM models have been employed in the industry for almost 30 years to identify and segment the customer base in countless companies across industries based on following questions: *How recently? How often? How much value?*. In our case we rely on the same RFM components to make the value modelling for topic impact dependent on the time and on the number of interacting users. Each component consists of a set of metrics based on the number of users interacting with the topic in a location, the engagement of these users with the topic -computed by the share of the content they produce related to the topic-, and their exposure to the content their network creates related to the topic, with the option of creating an aggregate index as well.

This paper is structured as follows: firstly the background information where we briefly review the related work is presented. Then, we introduce our method together with metrics to quantify the impact of a topic on the social media channel. After that, we present a system that implements our metrics and then we show some practical examples of topic impact quantification. Finally, we share our conclusions and point out future work on this topic.

2. Background and related work

In this section we provide all the background information and related work to base our research, starting with the review of impact modelling and topic diffusion, introducing the RFM model and finally discussing the approaches to topic modelling and information extraction in social media.

2.1. Topic diffusion and social media impact

The diffusion of news or topics in the social networks has been subject of intense research especially in the last years (Cavusoglu et al., 2010; Centola, 2010; Stieglitz and Dang-Xuan, 2013). Although the methodology we propose in this article is not intended to explain the dynamics of the topic propagation in the social networks, rather to provide a measure for the impact, there are common elements used in both researching lines to understand the contribution of a given user based on how active she/he is, the handling of the variation over the time of the topic-related activity and the semantic definition of the topic. Guille et al. (Guille and Hacid, 2012) defined three dimensions playing a role in the propagation of a topic: social, semantics and temporal to model the probability of dispersion. The social dimension is defined taking into account the users' activity index, the ratio of directed tweets to the user, the mentioning rate and whether the user being mentioned is directly related to the mentioner. On the other hand, the semantics is based on the presence of a keyword in the message being propagated. The temporal dimension is provided as a computation of the user activity in 6 partitions of the day, but probably leaving the door open to finer time granularity.

Rajyalakshmi et al (Rajyalakshmi et al., 2012) demonstrated the role of the strong links in the virality of the topics by modeling the diffusion with a stochastic approach, identifying as driving parameters the users activity time and the fading out effect –represented as a weight decay for a topic as time passed by—. In their work, two cases are clearly separated: users creating instances of a global topic or users copying it from their network –local social network effect vs. the overall trending effect—. Romero et al. (Romero et al., 2011b) established a mechanism relying on exposure curves to quantify the impact exposure to other users in making them adopt a new behavior (e.g.: turning them from passive to active contributors or to start using a hash-tag, etc). In addition, there have been several approaches to model the influence of a particular user in his/her own and in the global social media network. Ye and Wu (Ye and Wu, 2010) defined 3 different metrics to quantify the social influence: followers influence -the higher the number of followers, the higher the influence-, reply influence -the more replies one user receives, the more influential the user is- and re-tweet influence -the more re-tweets, the more influent-. Kwak (Kwak et al., 2010) suggested also 3 metrics but substituted the reply influence by one inspired by the Google Search PageRank algorithm (Page et al., 1998) to allow the propagation of influence. Depending on the metric applied the ranking of the top users varied. Romero et al (Romero et al., 2011a) demonstrated that influent users are those whose contributions are not just consumed but also forwarded and therefore overcome the so called passivity and more interestingly, that the popularity of an user and its influence don't quite often correlate. Cha (Cha et al., 2010) differentiated 3 kinds of influence for a social media user: due to the size of the user's audience or social network indegree influence-, due to the generated content with pass-along value retweet influence, which is also aligned with the passivity activity work presented by Romero (Romero et al., 2011b) and due to the engagement in others' conversation mention influence- and all of them are present as component for either Exposure or Engagement when applicable in our approach. The use of geo-localized SM interactions to provide information about local communities is a field of incipient research. In Scellato et al. (2011) the authors provide an extensive description of the social spatial properties of location based social networks. In Backstrom et al. (2010), the authors rely on the spatial proximity in combination with the social proximity to make geographical predictions. Another remarkable example (Cheng et al. (2011)) shows how to use location sharing services to explore and trace footprints.

2.2. RFM Background

The Recency, Frequency and Monetary (RFM) models were developed as a logical step in the evolution of marketing segmentation techniques. When the shotgun approaches (marketing everything to everybody) proved inefficient in terms of returns, the marketing campaigns started separating customers in segments based on socio-demographics attributes. Taking as segmentation criteria the customers' purchasing behavior proved more sensible indicator to response to campaigns than the former socio-demographic segmentation (Hughes, 2005). Especially the recency last time that a purchase was committed-, frequency how many purchases have been committed- and monetary value of the purchases committed- are usually employed to create a triple of scores per customer. The RFM models segment the customers relying on these scores, so that each segment is targeted in a particular much more tailored way. RFM approaches present also known limitations, like the risk of over-soliciting high-ranked customers, but this is rather a limitation related to the way of applying the findings of the model, not to the model itself, and therefore, it has no effect due to the way we want to apply it. Kumar et al (Kumar, 2008) pinpointed 3 limitations of RFM-based approaches to model customer behavior: it doesn't reveal any information about customers' loyalty which again is not an issue in our case, as the loyalty of an user to a topic is out of the scope of the metrics defined-, doesn't predict the next buy our model doesn't need to predict the next time the user is going to engage or be exposed to a topic- or the expected profitability over the time as our metrics work backwards, the predictive capabilities are not relevant-. Another metric traditionally related with RFM is the so called Customer Lifetime Value(Fader et al., 2005; Khajvand et al., 2011; Sohrabi and Khanlari, 2007) or the predicted value a customer is going to generate in her entire life time (Farris et al., 2010). We have not focused on Customer Lifetime Value like metrics as part of this study.

2.3. Social Media usage for topic extraction and trend detection

The social media platforms where users continuously post relevant messages referred to an ever changing huge variety of topics are the perfect playground for researchers to develop automatic topic uncovering algorithms. Blei et al. (Blei et al., 2003) started a new research line with their Latent Dirichlet Allocation, a multinomial probabilistic soft clustering of words based on co-occurrence. Many other researches have taken it as reference for topic extraction adding some improvements like in (AlSumait et al., 2009), where a method was suggested to prevent the generation of junk topics, or in (Chuang et al., 2012) it was pointed out the need for supervision by domain experts of the generated topics set, etc. Jones et al. (Jones, 1972) set the basis for the topic extraction based on the well-known Inverse Document Frequency (Salton and McGill, 1986). Latent semantic indexing (Deerwester et al., 1990) has also been vastly employed Automatic topics and trend detection have been also subject of research. Using Twitter particularly we found countless works oriented to extract topics and trends: for example, in (Cataldi et al., 2010) it was suggested an approach based on topic aging theory to extract the emerging topics only considering the authority of the user based on Page Rank; in (Naaman et al., 2011) the authors proposed a taxonomy of trends but specific to a given area relying on the extended geo-location capabilities of the social media conversation and suggested a trend featurization based on the associated messages to explain the different trends in local communities. In (Sakaki et al., 2010) the concept of *social sensor* was coined in order to detect events almost real time analyzing tweets. Mathioudakis et al. (Mathioudakis and Koudas, 2010) created a system called TwitterMonitor to detect and anticipate trends over the Twitter Stream.

3. Defining a new method for quantifying social media impact

In this section we present a new method and its metrics for quantifying how a topic impact the people of a place based on the RFM paradigm.

The motivation behind adopting the RFM approach in SM is two-folded. On one hand, it calls out and addresses separately each and every topic impact driver to latterly produce a compound metric or score: topics that are being discussed recently have a higher impact than topics that are no longer in the focus of attention -their relevance experiments a decay-; the repetition or how frequently a topic has been interacted with is also an indicator of impact, as well as the quality of the interaction with the topic -the value or monetary of the topic-. On the other hand, a broadly adopted, well-established approach that has already been in use for longer than a decade helps making the insights we create more actionable and directly usable for the industry. Unlike the traditional RFM approach, where recency is typically computed as how many *time units* back in the time (typically days) from the present moment (today/now), our approach is designed to be more generic: the time span is defined beforehand and the recency is always calculated taking as reference the extreme of the time span closer to the present. Additionally, and leveraging the emerging geo-localized nature of the social media interactions, our aim is creating a topic impact score for a specifying region, which unlocks a new set of possibilities –like impact comparison in two different cities or countries, etc–. Even if our metrics are design to overcome the existence of users with different levels of SM activity, there are several factors that might introduce certain bias, like the access to the internet of a particular region, the SM affinity –typically older population is less affine, etc– and other minor factors. Even for certain areas where the SM users constitute a less representative portion of the population, our model provides a result, but with higher volatility and less reliability, and should be interpreted as such.

3.1. Preliminary definitions

Before starting with the definition of our methodology, a set of concepts to support our metrics needs to be established:

Definition 1. *The set U represents the set of social media users from which we have evidence they have been in the location L (InLocation$(u_i, L, \triangle t)$) we are monitoring during the time period under analysis $\triangle t$*

$$U \equiv \{u\},\ \forall u_i \in U,\ InLocation(u_i, L, \triangle t) \tag{1}$$

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

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

$Follows(u_i, u_j)$ 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).

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

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

Definition 4. *We define all user interactions ($Interactions$) for a given user u_i over a time interval $\triangle t$, as:*

$$Interactions\ (u_i, \triangle t) \equiv \{it\},\ \forall it_i \in Interactions\ (u_i, \triangle t),\ Author(u_i, it_i, \triangle t) \tag{4}$$

A Social Media Interaction represents the atomic piece of content generated by the user u_i during the time $\triangle t$ in a Social Media Platform (e.g.: a tweet, a re-tweet). Thus, $Author(u_i, it_i, \triangle t)$ is a function that retrieves $True$ if u_i created the interaction it_i in the time period $\triangle t$, 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*.

We call *activeInteractions* to those made by any user $u_i \in U$ in the location and *passiveInteractions* the ones made by any user $u_j \in SN(U)$ the users in the location are exposed to.

Definition 5. *We define the set of Interactions for a given user u_i with the topic T over a time interval $\triangle t$ as:*

$$Interactions\ (u_i, T, \triangle t) \equiv \{it\},\ \forall it_i \in Interactions\ (u_i, \triangle t),\ Author(u_i, it_i, \triangle t)\ \wedge related(it_i, T) \tag{5}$$

Where $related(it_i, T)$ is a NLP membership function retrieving $True$ if the iteration it_i is connected to the topic T –intuitively, one or more words from the semantic field for the topic T are mentioned in it_i– and $False$ otherwise.

Definition 6. *We define Contributor to a topic T as the user u_i who created at least one interaction it_i with the topic T over a time interval $\triangle t$*

$$Contributor\ (u_i, T, \triangle t) \equiv True, \exists it_i,\ it_i \in Interactions\ (u_i, T, \triangle t), u_i \in U \cup SN(U) \tag{6}$$

Definition 7. *A user u_i is exposed to a topic T over the time span $\triangle t$, when there is at least one u_j in its social network $u_j \in SN(u_i)$ who contributed to the topic.*

$Exposed(u_i, T, \triangle t)$ is a logical function defined as:

$$Exposed\ (u_i, T, \triangle t) = \begin{cases} True, & \exists u_j, u_j \in SN(u_i), Contributor\ (u_j, T, \triangle t)\ is\ True \\ False, & otherwise \end{cases} \tag{7}$$

where $SN(u_i)$ represents all social network users connected to u_i

3.2. Recency in social media

In the traditional usage of the RFM model, *recency* has always been used as indicator for the last time a customer or prospect interacted with the brand, purchased a product, added a product to the online shopping cart, etc.*Recency* is traditionally used to assign a score to customers depending on how long ago the last interaction took place. In our case, we will provide the recency metric as an indicator to express how *up-to-date* the topic is. Thus, a topic that is *hot* today is going to have a much higher recency score than a topic the community stopped talking about weeks or months ago. As we are modelling a topic and not a particular user, we are going to suggest an aggregated approach. Depending on the topic in question, the interactions in the social media channel might be quite sparse, which introduces the need to work with thresholds. As pre-defined absolute thresholds might diminish the suitability for generic scenarios, we are going to set thresholds as a minimum of the topic share in a time unit (e.g.: a day), defined as follows:

$$Share(T, \triangle t) = \frac{\sum_{i=1}^{\#U}(\#Interactions(u_i, T, \triangle t))}{\sum_{i=1}^{\#U}(\#Interactions(u_i, \triangle t))} \quad (8)$$

Which is the sum of SM interactions created over the time period $\triangle t$ by the users geo-located in the place under analysis related to the topic T vs. the total number of SM interactions (those that are not related with the topic T as well).

For example, if we set the threshold to 0.1 per day, the *recency* would start counting up when the amount of posts related to the topic a day goes over 0.1 share of all interactions in the given day. This threshold can be adjusted to our convenience depending on the max number of different topics to be considered per day, the volatility tolerance, the topic in question, etc. The *Share* can be also defined in terms of Users engaged with a particular topic vs. all the active users in the particular day (which removes the bias introduced by overly active users, overly passive users, etc).

As explained before and unlike the traditional RFM model, our methodology is designed to work for any specific time frame and location in a more generic way. Hence, *Recency* does not have to necessarily be always measured from the present (today) back in the past. Rather, our definition requires the two extremes of a time interval $\triangle t$, so that both can be in the past. *Recency* is measured taking as reference the second extreme of the interval (*t_enddate*) over a time span to *t_startdate*. Although the most common time unit is the *day*, it's possible to adjust our approach to work in *weeks* or even in finer granularity units like *hours*, which is only advisable when the volume of interactions / time unit is sufficient to avoid volatility situations.

Based on the *share* concept to define *Thresholds* and the definition 6, we define *Recency* as follows:

$$Recency\ (T, \triangle t) = \begin{cases} \dfrac{t_i - t_{startdate}}{t_{enddate} - t_{startdate}}, & if\ \exists t_i, t_i \in [t_{startdate}, t_{enddate}] \\ 0, & otherwise \end{cases} \quad (9)$$

Where t_i is the first time unit closer to $t_{enddate}$, so that $Threshold \leq \sum u_i, Contributor(u_i, T, t_i) = True,\ u_i \in U \cup SN(U)$

3.3. Frequency in social media

In our method, *Frequency* is designed to measure how often interactions with the topic are registered during a defined time period. The more interactions with the topic, the higher the frequency and the higher the overall impact of the topic on the users located in the place under analysis.

Based on the type of interaction, we distinguish *frequency of exposure* or *passive frequency* and *frequency of engagement* or *active frequency*.

The *frequency of exposure* for a topic in a given period of time can be expressed as the number of users exposed to the topic per time unit. The subset of users exposed to the topic can then be defined as:

$$ExposedUsers(T, \triangle t) \equiv \{u\},\ \forall u_i,\ Exposed(u_i, T, \triangle t) = True,\ u_i \in U \quad (10)$$

Thus, the *frequency of exposure* can be defined as the number of users exposed over the time period:

$$Frequency\_Exposure(T, \triangle t) = \frac{\#ExposedUsers(T, \triangle t))}{length(\triangle t)} \tag{11}$$

Additionally, we define the *frequency of contribution* as the total number of users with active interactions with the topic in the specific time frame based on the set of all contributing users to the topic:

$$\begin{aligned} Frequency\_Contribution(T, \triangle t) = \frac{\#ContributingUsers(T, \triangle t)}{length(\triangle t)}, \forall u_i \in ContributingUsers(T, \triangle t), \\ Contributor(u_i, T, \triangle t) = True \end{aligned} \tag{12}$$

Putting both metrics together, we get to the envisioned *frequency* metric:

$$Frequency\ (T, \triangle t) = \frac{1}{2}(Frequency\_Exposure\ (T, \triangle t) + Frequency\_Contribution\ (T, \triangle t)) \tag{13}$$

In order to normalize these metrics, we make both relative to the quantification of the total exposure and the total contribution. For that we just put in relation the previously obtained frequency metric to the total number of users that could have been contributing or exposed:

$$Frequency\_Penetration\ (T, \triangle t) = \frac{1}{N}\ Frequency\ (T, \triangle t) \tag{14}$$

3.4. The Monetary component – Value in social media

Unlike *recency* or *frequency*, the *monetary* or *value* component requires certain modeling decisions about which factors need to be considered to which extent. When we talk about *value* for a SM topic, one could think of reach (Bogart, 1967). Measuring reach only, might leave certain aspects of the impact modeling unaddressed, like the *quality* of the audience where the topic is active, the level of engagement with the topic, the latent exposure among others.

Intuitively, *value* shall measure the number of SM users impacted by the topic -be it passively or actively-, quantify the intensity of this impact and put it in relation to the set of total users that could have been impacted. In the subsequent sections we are going to present our approach to model the different facets of a topic value to latterly consolidate everything into a single combined metric.

3.4.1. Modeling the exposure / engagement of a particular user with a topic

In our methodology we consider the level of exposure to the topic (concept inspired by the group contagion theory proposed in (Barsade, 2002)) and the level of user engagement with the SM content related to the particular topic as the main components for modeling value. *Exposure* builds upon the definition 7 and includes all scenarios where a given user could potentially read SM content related to a topic. As soon as the user adopts an active role towards the topic (creates or forwards content related with the topic), the user becomes a *contributor* (see definition 6) and we speak of *Engagement*. The table 1 presents the different categories we are going to use to attribute different intensity levels for both concepts.

Exposure Categories	
Disconnected Exposure	A social media user is exposed to the interactions of his/her social network (or further users directly linked to her). As the social media world doesn't stop but users regularly disconnect, there is so much content created within a particular social network that doesn't even get ever read by the user if online. With the proper set of web analytics in place, one could determine whether the user actually clicked on a piece of content generated in her social network or even model the probability of having read the piece of content based on session start time and session duration.
Connected Exposure	When there are interactions related to the topic within the user's particular social media network and the user himself/herself was online within the time window around the time of interaction with the topic. Active or online can be understood as connected and/or interacting with the social media platform (creating content, comments, etc).
Explicitly mentioned	User mentioned in a post related to the topic created by another. Unlike the previous categories of exposure based on the broadcasting of a message in the user's social network, this kind refers to peer-to-peer delivery of a message in the social channel: from a particular user to another particular user yet keeping it accessible to the entire social network of both. Even if we can't talk of an activity directly triggered by the user with the topic, the fact that a different user within his/her social network posted a piece of content about the topic and mentioned him/her on it, is going to increase the possibility of reading the post: (example from Twitter: user_2 posted: *"sorry mate, your team didn't have any chance against #manu @user_1"*. The user_1 gets a notification which most likely makes her reading the post from user_2 talking about a football match)
Engagement categories	
Active response	The user actually answered or commented a post created by another user within his/her social media network about the topic. (e.g.: based on Twitter: user_2 posted: *"sorry mate, your team didn't have any chance against #manu @user_1"*. user_1 replied *"@user_2 Chelsea FC for ever! #cfc"*)
Active forwarding	The user just confirms that he/she feels identified with a piece of content generated by another user within her social media network about the topic we are analyzing (depending on the social media platform as a "I like" or as a *Retweet*).
Actively initiated	The user starts talking about a topic within his/her social network. She is the initiator and the one who brought up the topic into her social network. We see this one as the highest level of engagement.

Table 1: Exposure and Engagement categories

To enable finer granularity when we measure the intensity within the *Engagement* categories, we additionally use the type of content. For example, one user taking a picture and posting it with a message about the weather shows more engagement that just a text. The same applies for links: a user sharing a link about a topic suggests that the user read about the topic already somewhere else and shared the reference to this content indicating a higher level of engagement as well.

To model the degree of engagement and exposure related to a topic based on the categories defined in table 1 and including the different content types, we apply a weighting schema. To apply the weights we partition the set of all active interactions of a given user u_i with the topic T in the time interval $\triangle t$ by the type of content C on one hand and by Engagement Category En on the other hand: The types of content we considered $\{Text\ only$, $Contains\ links$, $Contains\ video\ or\ picture\} \in C$ constitute complete partitions of $Interactions\ (u_i, T, \triangle t)$, so that $\forall it_i \in Interactions\ (u_i, T, \triangle t)$, $\exists!\, c_j, ContentType(it_k) = c_j, c_j \in C$. And so do the Engagement categories $\{Active\ Response$, $Forwarding$, $Active\ Initiating\} \in En$, so that $\forall it_i \in Interactions\ (u_i, T, \triangle t)$, $\exists!\, e_j, EngagementCategory(it_k) = e_j, e_j \in En$

Let's express the different partitions of the Interactions set for a given user based on types of content as

$Interactions\ (u_i, T, \triangle t \mid c_k), c_k \in C$ and the partitions based on engagement categories as $Interactions\ (u_i, T, \triangle t \mid e_k), e_k \in En$.

Based on both partitions, we define *Engagement* by weighting the different *engagement category* - *content type* pairs, as follows:

$$Engagement(u_i, T, \triangle t) = \sum_{j=1}^{\#(En)} \sum_{k=1}^{\#(C)} w(c_k, e_j)\ \#\ ((Interactions\ (u_i, T, \triangle t \mid c_k)\ \cap\ Interactions\ (u_i, T, \triangle t \mid e_j)) \tag{15}$$

Similarly, the set of all Interactions created by the Social Network of a given user u_i, $SN(u_i)$ related to the topic T in a period of time $\triangle t$ can be partitioned based on Exposure categories for the user u_i $\{DisconnectedExposure$, $ConnectedExposure$, $ActivelyMentioned\} \in Ex$, so that $\forall it_k \in Interactions\ (SN(u_i), T, \triangle t)$,
$\exists!\, e_j, ExposureCategory(it_k) = e_j, e_j \in Ex$. Ex is the set of all partitions based on exposure or exposure categories. Based on these partition, we can define *Exposure* as follows:

$$Exposure\ (u_i, T, \triangle t) = \sum_{k=1}^{\#(SN(u_i))} \sum_{j=1}^{\#(Ex)} w(e_j)\ \#\ (Interactions\ (u_k, T, \triangle t)\ \cap\ Interactions\ (u_k, T, \triangle t \mid e_j)) \tag{16}$$

3.4.2. Active and Passive Impact based on Engagement and Exposure

The metrics previously defined to quantify *Engagement* and *Exposure* work on absolute terms. The definition of Impact upon them puts both metrics into relation to the total number of interactions created by the user or generated within the SN of the user. To address that, we define *Active Impact* or *Engagement Index* and *Passive Impact* or *Exposure Index* as follows:

$$ActiveImpact\ (u_i, T, \triangle t) = \frac{Engagement\ (u_i, T, \triangle t)}{\#Interactions(u_i, \triangle t)} \tag{17}$$

$$PassiveImpact\ (u_i, T, \triangle t) = \frac{Exposure\ (u_i, T, \triangle t)}{\#Interactions(SN(u_i), \triangle t)} \tag{18}$$

The resulting value component of the impact of a topic T on a given user u_i over a period of time t is a combination of *Passive Impact* and *Active Impact*:

$$ImpactValue\ (u_i, T, \triangle t) = \frac{1}{k + l}\ (k\ *\ ActiveImpact\ (u_i, T, \triangle t) + l\ *\ PassiveImpact\ (u_i, T, \triangle t)) \tag{19}$$

where k represents the weighting for the Active Impact and l represents the weighting for the Passive Impact (l and k are positive numbers). Depending on how the use case gives priority to the engagement over exposure, the values for the weights l and k are defined.

This is the definition we suggest for the *Monetary* or *Value* component in our RFM methodology.

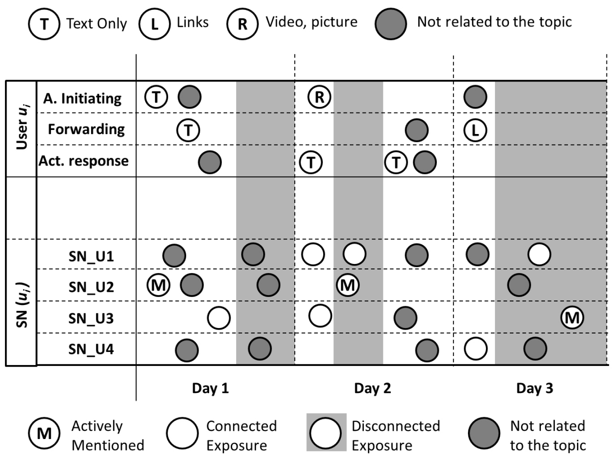

Engagement

		Actively Initiating	Forwarding	Active Response
Weights	Rich Content	1	0,8	0,6
	Links	0,8	0,64	0,48
	Text Only	0,7	0,56	0,42
# Topic Interactions	Rich Content	1		
	Links		1	
	Text Only	1	1	2
		Engagement		3,74
		Total Interactions		11
		Engagement Index		0,34

Exposure

Weights	Disconnected	0,2
	Connected	0,5
	Antively Mentioned	1
# Topic Interactions	Disconnected	2
	Connected	4
	Antively Mentioned	3
	Exposure	5,4
	Total Interactions	21
	Exposure Index	0,257

Figure 1: Example of timeline with all interactions of user u_i and his/her Social Network $SN(u_i)$

Figure 1 explains how both Engagement and Exposure Indexes are obtained for a fictive user's time-line over 3 days applying the formulas (17) and (18). If we weighted both components in 19 with 0.5 each, the total impact of the Topic on the user u_i would be 0.298.

The formula (19), which defines the Impact at user level, can be extended to all the users U in the location:

$$ImpactValue\,(U, T, \triangle t) = \frac{1}{\#U} \sum_{i=1}^{\#U} ImpactValue\,(u_i, T, \triangle t) \tag{20}$$

Building on top of (Romero et al., 2011b), our impact modeling assigns different engagement levels depending on whether the user initiates the topic within her social network or just engages with a topic currently discussed in her network. Unlike (Kwak et al., 2010), our approach just considers the inbound diffusion component, what we called exposure but omitting the outbound diffusion —the geo-located users are not necessarily exposed to the activity of their followers according to the way the online social media platforms are designed on one hand and the engagement of the followers of geo-located users doesn't contribute to the overall engagement in the location under analysis on the other hand—. The Exposure and Engagement components we suggest relies on the activity / passivity concept introduced in (Romero et al., 2011a), evolving into the set of metrics we defined above. As we mentioned before, we've chosen exposure and engagement to model our value metric because both represent a quantification of the topic intensity on the users located on a particular location.

3.5. *Quantifying the topic impact using the Social Media RFM model*

Our social media RFM model provides a single value quantifying the impact of a topic in a SM channel within a period of $\triangle t$ as a combination of the metrics discussed above:

$$Impact\ (U, T, \triangle t) = F\ (Recency\ (U, T, \triangle t),\ Frequency\ (U, T, \triangle t), Value\ (U, T, \triangle t)) \tag{21}$$

The function F can be any combination of the three components, varying from a simple average to a weighted average to more complex scenarios. The definition of F should be made dependent on the use cases (i.e.: the recency component might be ignored or weighted very low for recurring topics like "weather" under the assumption that every body post about the weather regularly, whereas other scenarios like a particular event like "New York Marathon"). As aforementioned, having a single index, allows for combining the impact of different topics, in different locations over different periods of time. For example, you could compare the topic "Pope election in Rome during the conclave weeks" with "Royal wedding of Prince William the week of the 29th of April 2011 in London".

4. System Architecture

Even if the metrics presented so far can be applied to each and every social media platform — provided there are means of getting access to the required information—, the system we implemented focuses on Twitter only. We've chosen Twitter over other existing networks because of the ease of information extraction (no constraints in terms of the need for being connected to users to retrieve them over the APO), because of the variety of the topics discussed unlike other purpose specific SM networks –like Linked.in–, because it's broadly adopted, and because the geolocation capabilities are extensively developed.

The system polls the geo-located tweets from the publicly available Twitter Search API [2] , flags those tweets that are related to the topic, extracts the information about the users involved in these tweets, including their social network and finally applies the set of metrics for impact modelling. The system consists of 4 different modules in charge of different labours all along the process.

Each module consists of a set of components with a clearly defined function. In the following sections we are going to describe how the different modules work and what the role of the components being involved is.

4.1. *Tweets harvester*

This module performs poll-requests from the Twitter Search API to store the tweets into a local data base for further processing. The tweets are selectively picked for a given area which is configured in the harvester, namely

[2]Available at `https://dev.twitter.com/docs/api/1/get/search`

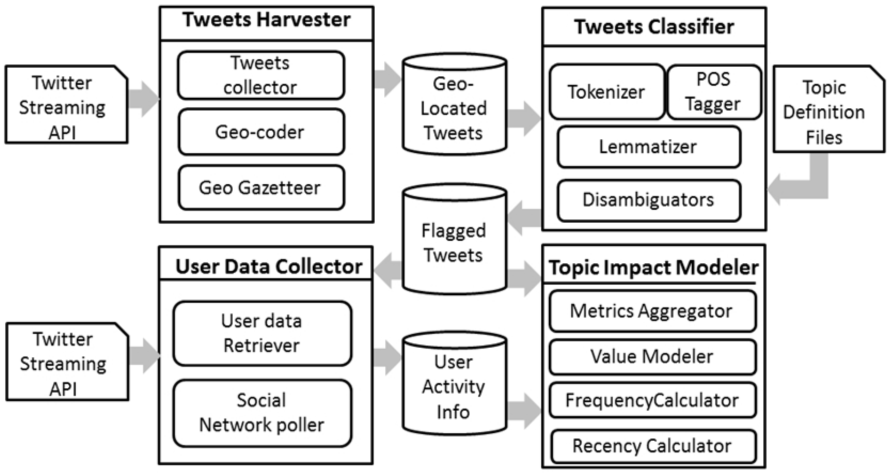

Figure 2: System Modules Overview

the one we want to perform the topic impact analysis over time. Additionally, the Twitter API supports the filtering by language (e.g.: only tweets in English), but even it would make the later NLP much easier, it might disregard the interactions of all users related with the topic in the target area for being in a foreign language. We opted for a work-around that doesn't filter out the tweets by language upfront, yet doesn't introduce the need for applying NLP techniques in all identified languages, as we are going to explain in the next section.

The harvester also provides the capability of assigning the interactions to already standard geographical output systems (like postal sectors, output areas, census units, etc depending on the country).

A gazetteer supports both correction of inaccurate information and spelling mistakes in addresses, etc. To implement these functionalities our system relies on existing geo-coding API's provided by the major web mapping providers[3] -which usually are free of use up to a limit of request per day-. The geo-mapper component takes as input the shape files (polygon lines) describing the output geographical units of a system of choice (e.g. postal sectors) and applying the standard point-in-polygon algorithm (Sutherland et al., 1974) establishes the mapping of the interaction or tweet to a standard geographical unit. The outcome of the harvester is a collection of full-fledge[4] tweets with a time stamp, a pair of geographical coordinates and potentially the link to a particular standard geography unit.

4.2. Tweets classifier

The mission of this module is basically separating all the harvested tweets that are related to the topic from the others.

[3]Google Geocoding API: `https://developers.google.com/maps/documentation/geocoding/` The location API provided by Bing Maps REST Services: `http://msdn.microsoft.com/en-us/library/ff701715.aspx`

[4]Full-fledge because all the meta information coming from the Twitter API has not been discarded (see `https://dev.twitter.com/docs/platform-objects/tweets`)

Running the system to produce the impact metrics for a given topic requires the non-trivial task of gathering and structuring the set of keywords that qualifies the topic. We suggest following sources:

- Social Media Entities related to the topic: Set of official accounts, nicknames, hashed tags, etc. users mention in their interactions with the topic (e.g.: for the topic "tennis", we would have RafaelNadal for Rafa Nadal, DjokerNole for Novak Djokovich, etc). For completeness it should include both official accounts and those that are not official but with high levels of activity.
- Topic Named Entities: set of named entities related to the topic (e.g.: "Rafael Nadal", "Noval Djokovic", etc)
- Topic Lexicon File: containing the set of non-named entities related to the topic (e.g. in the *tennis* domain: "ace","match ball", "set", "advantage", etc.)

The orchestration of the steps required to perform the Tweets classification is designed to minimize the number of *false positives* as soon as possible in the process and increase the overall performance of the classifier. Thus, the number of tweets remaining from one step to the subsequent one is intended to shrink. Likewise, the steps involving higher complexity are pushed towards the end, whereas the simple and efficient ones are done at the beginning.

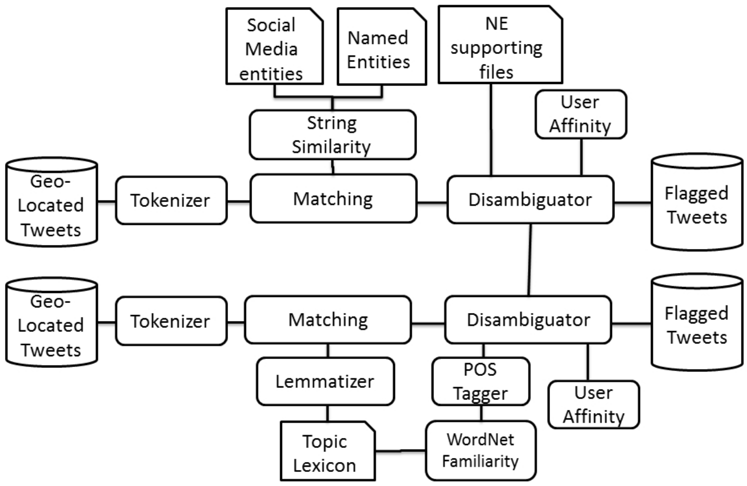

Figure 3: Double-pipe Tweets Classifier

The way the classifier works in 2 phases, is explained in Fig. 3 both pipes share the same tokenizer: each geo-located tweet is tokenized applying a sentence tokenizer first and a word tokenizer later (based on (O'Connor et al., 2010)) both adapting the Punkt Tokenizer (Kiss and Strunk, 2006) to deal with social media texts. The modified tokenizer provides the stop words removal as well. The Social Media and Named Entities pipe intents to match each and every reference term listed in the social media and Named Entities files applying a string similarity algorithm (Yang et al., 2003), which delivers a similarity score. The matching module in our system then implements thresholds –which differs depending on the source— to support the fact that the social media content is often full of spelling errors, which is likely to happen even more frequently when it comes to named entities of foreign people (e.g.: staying in tennis, Nalbandian is often spelled as Nabandian even by renowned tennis twitter accounts) (Clark et al., 2010). Even if our unigram based approach might look simple compared with more sophisticated approaches like hierarchical Dirichlet bigram language models (MacKay and Peto, 1995) or based on semantic *gists* (Griffiths et al., 2007), the suggested approach for the disambiguation allows for keeping the topic model simple and therefore less processing intensive.

For all the messages tagged positively as belonging to the topic and added the set of matching terms or keywords, a disambiguation process takes place. The disambiguation relies on a semi-automated supporting list of homonyms for the named entities:

1. When the term to disambiguate has not been the only one part of the tags for a tweet and one or more of the other tags was univocally related to the topic, the context was sufficient for the disambiguation (E.g. based on topic "football" *'RT bbc5live: 12:45: Our 1st #BPL commentary of day- Newcastle v Liverpool...a point will take Brendan Rodgers side top #NUFC #LFC'*, tagged with 'lfc,nufc,lfc,liverpool,rodgers' - *Rodgers* is disambiguated by the presence of other tags related to the topic).
2. When just a single term to disambiguate is part of the tagging set for a tweet and this term is a Named Entity, we applied a technique based on the expansion of the named entities by related terms like surname, name, alias, etc. inspired by (Fujita and Fujino, 2013). E.g.: *'@Theleaguemag a young steve bruce !!'*, *Bruce* disambiguated by 'Steve Bruce'
3. When the disambiguation is not possible as none of the points mentioned above can be applied, the term is marked for a new run of the disambiguation once the pipeline for the topic lexicon is done and also matching lexicon terms might have been added to the tweet.
4. If no disambiguation is possible based on the tweet itself, we try to leverage the affinity of the tweets from the same author with the topic. If the affinity is high, the chance of the tweet to be related to topic is higher, even if there's some room left to interpretation and we might have to accept certain tolerance. The affinity of the user to the topic is calculated as a ratio of the positive tweets over the total number of tweets that have been harvested for the user.

The topic lexicon pipe works likewise in 2 phases, the matching phase and the disambiguation phase. As input for the disambiguation, the associated polysemy coefficient is calculated using the so called WordNet familiarity (Miller, 1995). Each term is basically given the number of synsets, which helps us understanding when a disambiguation is required. To disambiguate lexicon terms, we apply following approaches:

1. Part of Speech based rules: after the POS tagging, for which we rely on the tailored tagger for social media (e.g.: *'was reading about arthritis drugs and apparently sometimes your hair falls out what if this happens to me'*, 'readingfc,reading' (*reading* - VBG - Verb, gerund/present participle- disambiguated to non-related to

the topic) Reading . *'Reading FC is gonna have a hard year'* -NNP - Proper singular noun- disambiguated as related to the topic 'football').

2. Presence of multiple terms of the terms reference set (both Entities and other lexicon terms).
3. Author's affinity to the topic (as explained before).
4. Supervised disambiguation for the most frequently identified stand-alone terms.

Even if our approach to disambiguation is not proven to work in 100% of the cases, the number of false positives is to certain extent balanced by the terms that have not been considered in the modelling topic sources. Gathering all the terms that could identify a topic is hard task whose complexity increases depending on how dynamic topics are.

4.3. User Data Collector

This component is in charge of polling the Social Network of all the authors of the tweets identified by the first pass of the classifier (*socialNetworkUsersSet*). Additionally, the harvester comes again into picture to retrieve the tweets of each and every user in each author's social network to enable the exposure analysis. Once all tweets of all users in the social network for the time-period being analyzed have been stored, the classifier acts again to flag those belonging to the topic (*socialNetworkTweetsSet*).

Another important task carried out by this component is the implementation of the exposure window based on each user's interaction and each user's social network activity index.

4.4. Topic impact modeler

After gathering all the social media content and classifying it according to how related to the topic under analysis it is, this module applies the metrics computing the exposure, the engagement and recency.

This module implements a pre-processing stage consisting of following steps:

- Contributors' flagging: setting the contributors' flag for all the harvested users, if they authored any of the tweets flagged as related to the topic (as defined in (6)).

- Exposed Users flagging: considering the *socialNetworkUsersSet* and taking as input the outcome of the previous step to determine who in any user's social network is a contributor, as defined in (7).

- Content categorization: classification of the tweets flagged by the Tweets Classifier and including the *socialNetworkTweetsSet* based on Exposure and Engagement categories explained in the sub-section 3.4.2.

Applying the equation (8), the *Recency Calculator* computes the share of the topic based on the ratio of flagged Tweets vs. all harvested Tweets, which is used as a reference only. The share is going to provide an indication of how reliable and how significant the metrics are. The share is also calculated in the particular time units specified (e.g.: day) backwards from the *t_enddate* until we reach a time unit whose share is greater than a threshold. The threshold selection in particular but also the need for the Recency component in general, depend very much on the specific use case. The Recency value is calculated according to the equation (9).

The *Frequency Calculator* then takes over to compute both the Frequency Exposure as defined in the equation (11) relying on the pre-computing stage results for Exposed Users and the Frequency Contribution, like in the formula (12). The results are then combined into the overall frequency calculation as defined in the equation (13). The last step accomplished by this component is normalizing the result as described in the formula (14) to provide the Frequency Penetration value.

The *Value Modeler* is in charge of producing the value component of the framework. This component computes first the topic engagement or active impact as defined in (17) and the passive impact or topic exposure (18) later. Both computations require the previous content categorization we described as a pre-processing step above.

As a final step, the *Metrics Aggregator* pulls all the metrics together to generate a single value applying a concrete implementation with a particular set of weights per component as defined in the function (21), if a general score is desired.

5. Evaluating the Social Media RFM Model

The purpose of this section is to evaluate how the suggested RFM model behaves in different scenarios to prove both sensibility and usage of the set of defined metrics.

Our proposal to the validation consists of two real world topics to ensure the coverage of all the variety of events a topic might manifest. For each one of the suggested topics we are going to present different scenarios carefully selected to thoroughly demonstrate the performance of our framework while highlighting the role of each metric in the different cases.

Our first topic revolves around two events that shook the hearts of multitudes within the space of one week: the deaths of the famous American actor *Paul Walker* [5] on November the 30th 2013 and the decease of the charismatic Peace Nobel Price winner, South Africa first black president and anti-apartheid icon *Nelson Mandela* [6], just 6 days later. These one-off events are going to help us demonstrate the role of recency and frequency taking different time analysis windows (centered on the day of the death, the week after, the week before, etc). Additionally, we are going to use provide an impact comparison of both deaths in the considered locations.

As second topic we chose *Football*, much wider in scope but highly suitable to prove the performance of our metrics due to the following reasons: recurrence –there are regular matches coming every week–, variety of scenarios –team playing at home, as visitor, national championship, Champions League, etc.–, fine granularity in time –which allows for engagement and exposure calculations for hourly intervals–, popularity –with a lot of social media content generated about the topic and therefore, less volatile– and easy to model –with a rather large volume of SM and

[5] `http://www.bbc.com/news/world-us-canada-25173331`
[6] `http://www.bbc.com/news/world-africa-25249520`

Named Entities and comparatively small lexicon, which natively reduces the need for disambiguation and therefore, the number of false positives–. When the complexity of the topic increases, for example due to the presence of more ambiguous terms, the number of terms required to ensure a proper coverage, the need for remodelling for rapid changing topics, etc., the quality of the topic model might be affected and consequently the quality of the metrics we suggest in this paper. Depending on the complexity drivers, more advance modelling techniques could be applied to tackle particular problems, like disambiguation, etc. but they escape the scope of this paper.

Figure 4: Hourly distribution of Tweets gathered on Dec. 1st by the 1km and 5km harvesters

5.1. The set up

We carried out the set up of our evaluation in three different steps: gathering of all the required information to create the topic definition files for the topics, configuration of the harvesters to start gathering geo-localized tweets in the locations of interest and definition of the impact aggregation function and weighting schemas.

For Mandela's and Paul Walker's deaths we just took the named entities of both personalities and the popular aliases people use to refer to them (e.g.: *Madiba* for Nelson Mandela).

The football topic definition file, due to its breath, has not been that straight away. We gathered the named entities from the official site of the Premiere League[7] , the official twitter accounts from all the players[8] and teams[9]. The lexicon file has been manually created compiling several sources to extract the football specific terms (78 unique terms).

Table 2 shows the summary of the sources we've employed to characterize the topic we are currently analyzing: 393 players', managers' and clubs' official accounts, 52 hash-tags representing football clubs in the group *Social Media Entities*, 692 names of players, managers, clubs (*Named Entities*) and 72 additional terms related to football (*penalty, offsider, etc*).

[7] `http://www.premierleague.com/en-gb/clubs.html`
[8] `http://www.transfermarkt.co.uk`
[9] `http://footballersontwitter.com`

Topic Characterization: Football In UK

Type	Entity group	#Terms	Type	Entity group	#Terms
Social Media Entities	Club Official Accounts	49	*Named Entities*	Players	570
	Players Official Accounts	315		Teams	44
	Managers Official Accounts	29		Managers	29
	Club Official Hash-tags	52		Clubs	49
Lexicon	Football terminology	72			

Table 2: Topic Characterization: Football In UK

We set up 5 harvesting engines: 2 of them to monitor the activity on two well-known football stadiums: Stamford Bridge (Chelsea FC) and Old Trafford (Manchester United FC), 2 additional ones centered on both stadiums but with a much larger radius (5km) covering an important part of London and Manchester and a last engine also with a radius of 5 km covering the city of Edinburgh (a place a priori not so much related with the topic *football*).

Although our harvesters have been running for longer than 3 months, we are going to focus our analysis on the first two weeks of December 2013, where the vast majority of scenarios manifest. The harvesters gathered 1088627 tweets during these 2 weeks in the mentioned locations.

As example of harvesting we show in Fig. 4 the geographical distribution of the tweets in Manchester for different hours for the 1st of December.

In order to make results comparable across all topics' analysis, we applied the same weighting schema for the Exposure Groups and Engagement Categories and Content Types. The definition of the weights is the one used in Figure 1 to explain how the metrics are calculated.

The weightings to compute the Impact aggregating both Active and Passive components –see equation 19– are going to be 0.5 for each one. Even if our suggested metric provides the flexibility of making a component prevail over the other one by increasing its weight, the examples we are presenting here don't require any special handling. For consistency reasons –especially to make the outcome comparable– we apply the same weighting schema to all examples.

The Aggregation function when we really require a single score (see equation 21) is going to be defined using the Frequency Penetration as Frequency component and weighting the value component twice as much as the other two: 10 for Recency, 30 for Frequency and 60 for Value.

5.2. *Topic 1: Impact modeling for Mandela's and Paul Walker's deaths*

Paul Walker died in a car accident the 30th of November approximately at 23:30 GMT. Nelson Mandela died on the 5th of December short before 19:00 GMT after being hospitalized. The SM platforms echoed both events and our harvested gathered the users' reactions in all 5 locations.

As aforementioned, both events share the same pattern: one-off, high SM resonance, a very quick ramp-up phase to a peak and a rather short fade-out phase to practically disappear after a few days.

Fig. 5a) shows the active impact (as defined in the equation 17) over all users identified per location. The small-radius harvester in Chelsea shows a higher score than the others in the first day but stabilizes one day after. The passive impact though, as displayed in Fig. 5b) positions Edinburgh and the 5km Chelsea harvester, specially the day after the decease, much higher than the small-radius ones. The less exposure of Manchester 1 km users is remarkable. Fig. 6a) shows the daily aggregated view of the value, which is aligned with the results of the previous charts: Chelsea and Edinburgh more impacted by the death than Manchester.

Manchester is on the other hand, where Paul Walker's death caused a higher impact than in Edinburgh –see Fig. 8 and 9a)–. The reason might be related to the different affinities of the Edinburghian and Mancunian local tweeting communities. Again, this is one example of the capabilities of our framework to understand localized preferences and people profiles. The long tail values in both cases show a very slight increase 4 days after the deceases, which might be triggered by TV media showing *TV specials* about these personalities and/or best movies in the case of Walker.

Fig. 7a) and Fig. 10a) provide the value overview by hour for both deaths. We observe a similar pattern: all monitors showing a peak which lasts for 4-5 hours to then go down. The hourly impact score right after the announcement of the tragic event topped 0.4, which is comparable with the highest impact values reached during football games in their local stadia (see in Fig. 12 the Manchester harvesters when the Manchester United - Everton match was kicked off at 07:45 pm in Old Trafford).

Figs. 7 and 10 present the daily frequency for both events, each one with a suggested threshold for the Recency calculation. The Table 3 shows the value for the different components for one week. Even if the value component taken the death's day and the day after is similar for both Mandela and Walker, the fact that we are considering the entire week (1st to 7th Dec.) and Walker's death took place right at the beginning, make the Recency substantially differ in favor of Mandela's impact, which we see in the resulting score. The charts in Fig. 13 help understanding the hourly value distribution in both cases over the week per harvester.

In Fig. 14a) we've taken the hourly impact value obtained for both topics the day after both personalities passed away. We show the difference between both values per hour and per location, to demonstrate how making the impact quantifiable enables the comparison. Thus, we can see that in the locations under analysis, Mandela's death had a greater impact during the first two to three hours, whereas Walker's death outperformed it in the next 4 hours especially in the greater Chelsea area. Both impact values get closer as the day proceeds.

5.3. Topic 2: Impact modeling for the topic Football

The first two weeks of December have been very intense in terms of football matches for both Chelsea FC and Manchester United FC, as we can see in Table 5.

Fig. 11a) shows very well the effect of a club playing in its home stadium in terms of impact. On Sunday the 1st, the active impact close to Old Trafford –harvester MANU 1 km– is very low compared with next Wednesday and Saturday. The same behaviour is shown by the Chelsea 1 km harvester: high impact on the Sunday and pretty low on the next Wednesday and Saturday, which again comply with the fixtures given in the Table 5.

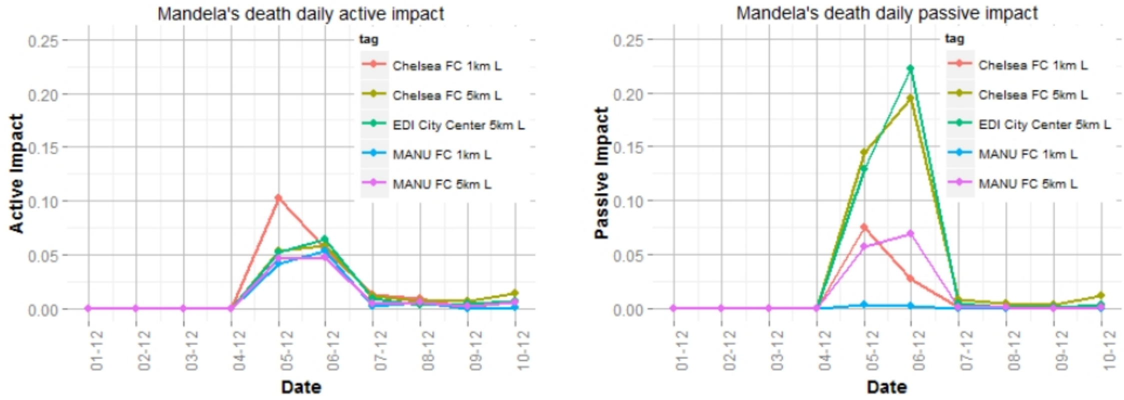

Figure 5: Mandela's death active (a) and passive(b) impact

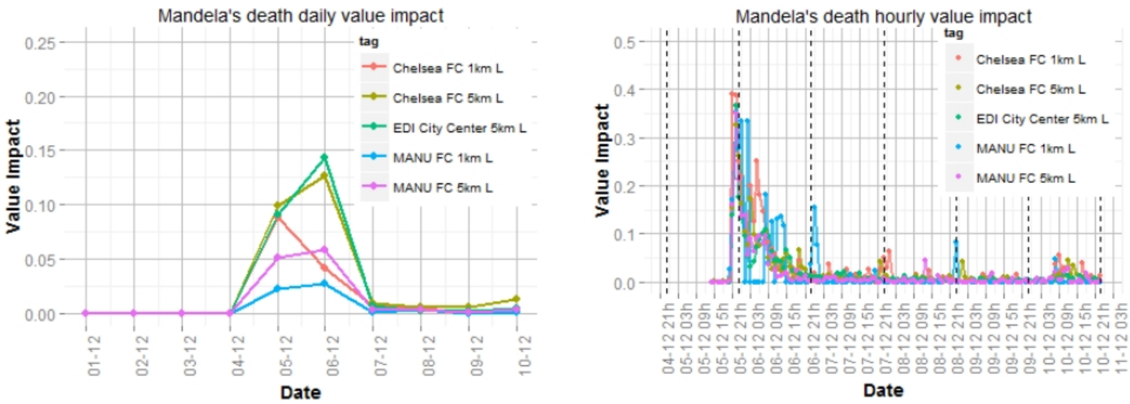

Figure 6: Mandela's death daily (a) and hourly (b) impact value

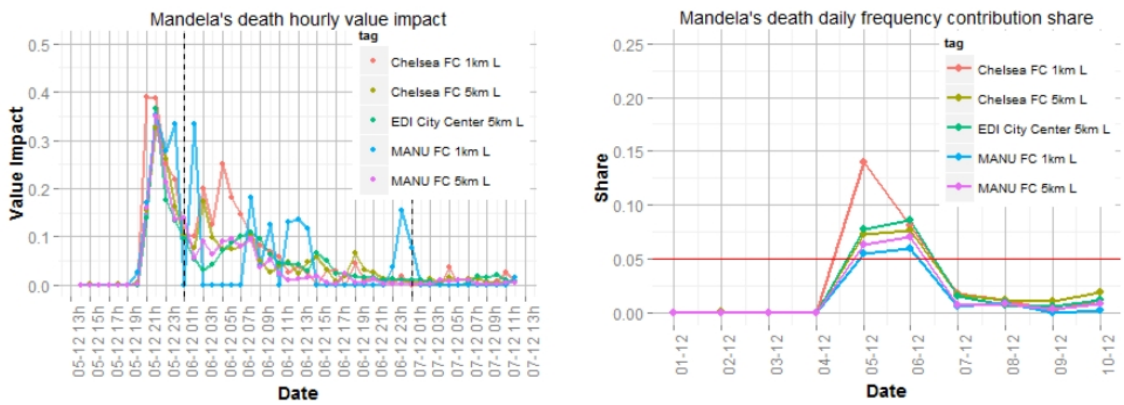

Figure 7: Mandela's death hourly detailed impact value (a) and frequency (b)

Walker's vs. Mandela's death Impact 1st-7th Dec. 2013

Topic	Harvester	Recency	Frequency	Value Impact	Resulting Impact
Mandela's death	*Chelsea 1km*	0.85	0.04	0.028	11.41
	Chelsea 5km	0.85	0.02	0.039	11.44
	MANU 1km	0.85	0.021	0.013	9.95
	MANU 5km	0.85	0.021	0.021	10.40
	EDI 5km	0.85	0.03	0.042	11.92
Walker's death	*Chelsea 1km*	0.14	0.018	0.01	2.89
	Chelsea 5km	0.14	0.033	0.035	4.54
	MANU 1km	0.14	0.024	0.015	3.02
	MANU 5km	0.14	0.034	0.025	3.97
	EDI 5km	0.14	0.021	0.018	3.15

Table 3: Impact modeling of the Mandela's and Walker's deaths on 5 locations

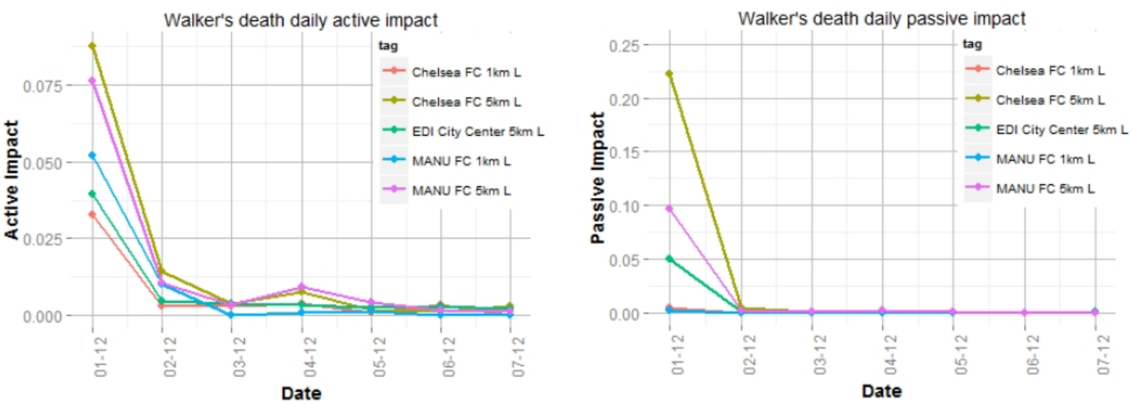

Figure 8: Paul Walker's death active (a) and passive(b) impact

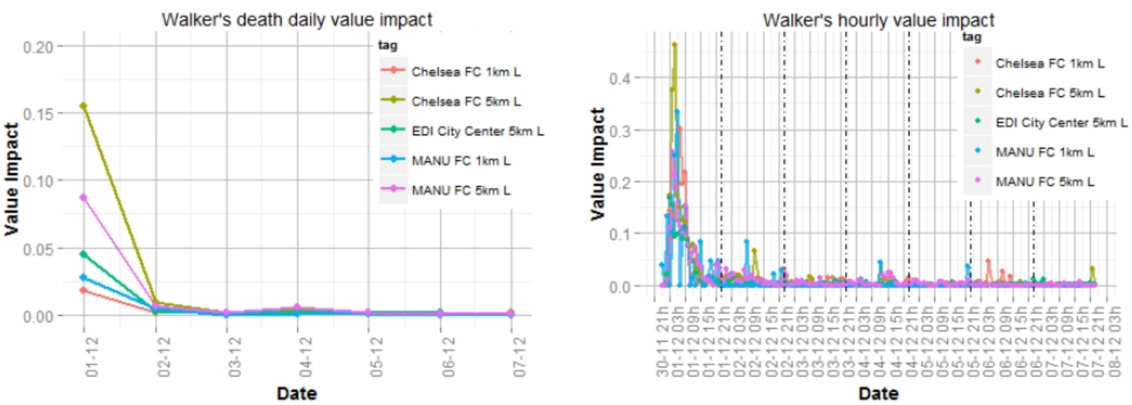

Figure 9: Paul Walker's death daily (a) and hourly (b) impact value

When we increase the radius and move away from the stadia, the differences between the impact when the local

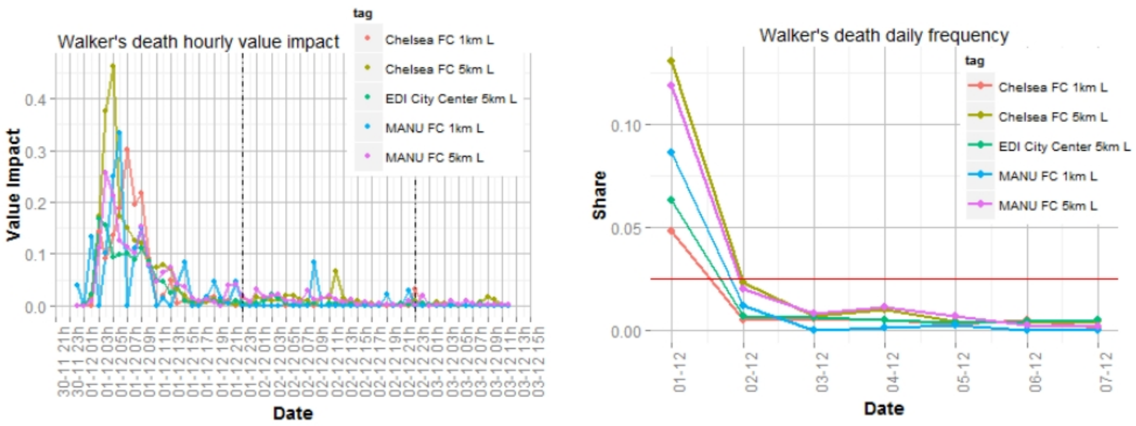

Figure 10: Paul Walker's death hourly detailed impact value (a) and frequency (b)

club plays at home or as visitor expectedly diminishes. It applies to both active and passive impact and to the aggregation of both (Fig. 11) . It's remarkable the low impact of the topic football in the city of Edinburgh or just the low interest on the Premier league.

Fig. 14 b) shows the week vs. week impact comparison per harvester. The impact of the home matches is obvious, but also the importance of Champions League in the 5 km harvesters. Edinburgh stays "unimpacted" by the topic Football, as we said before.

Football Impact 1st-7th Dec. 2013

Topic	Harvester	Recency	Frequency	Value Impact	Resulting Impact
Football	*Chelsea 1km*	1	0.387	0.051	24.721
	Chelsea 5km	1	0.112	0.102	19.514
	MANU 1km	1	0.087	0.042	15.207
	MANU 5km	1	0.11	0.107	19.964
	EDI 5km	0	0.01	0.01	0.9

Table 4: Impact modeling of the topic Football on 5 locations

In the Table 4, we present the impact results for the topic football in the defined week. We see the Chelsea 1 km harvester leading the table with a score of over 24 points. MANU FC showed a comparably lower impact – almost 10 points back–. The 5 km MANU and Chelsea harvesters, not so sensible to when the local club plays in the local stadium, show remarkably similar results. Chelsea FC and MANU played 3 times each during the week, but Edinburgh stayed completely indifferent to that.

6. Conclusions

In this paper we present a new model built upon geo-localized Social Media interactions to quantify the impact of a topic on a particular location and to monitor how it changes over time. As a foundation for our model, weve

1st-14th Dec. 2013 Football Fixtures				
Date	Competition	Home	Result	Visitor
Sun 01.12.13	Premier	**Chelsea**	3 - 1	Southampton FC
Sun 01.12.13	Premier	Tottenham Hotspur	2 - 2	**Manchester United**
Wed 04.12.13	Premier	Sunderland	3 - 4	**Chelsea**
Wed 04.12.13	Premier	**Manchester United**	0 - 1	Everton
Sat 07.12.13	Premier	Stoke City	3 - 2	**Chelsea**
Sat 07.12.13	Premier	**Manchester United**	0 - 1	Newcastle United
Tue 10.12.13	Champions	**Manchester United**	1 - 0	Shakhtar Donetsk
Wed 11.12.13	Champions	**Chelsea**	1 - 0	Steaua Bucuresti
Sat 14.12.13	Premier	**Chelsea**	2 - 1	Crystal Palace FC

Table 5: Chelsea and Manchester United matches calendar

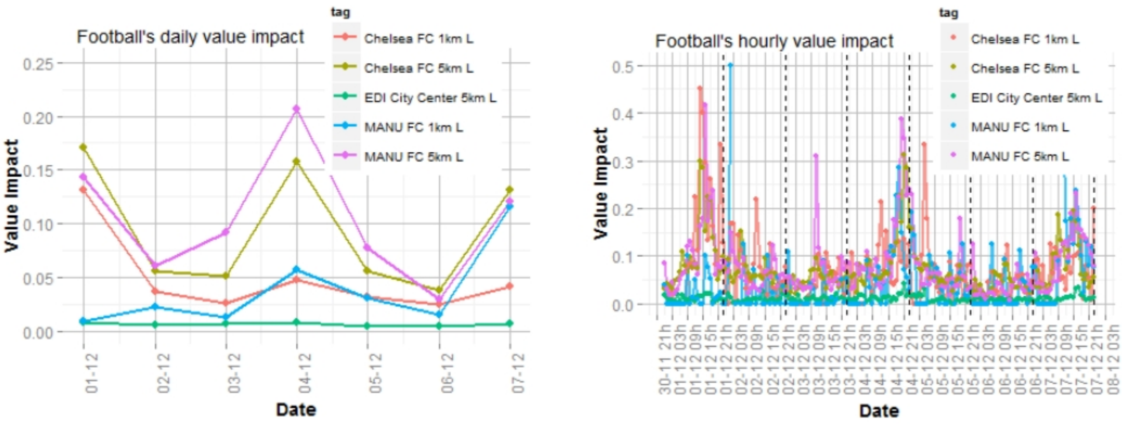

Figure 11: Football daily (a) and hourly (b) impact value

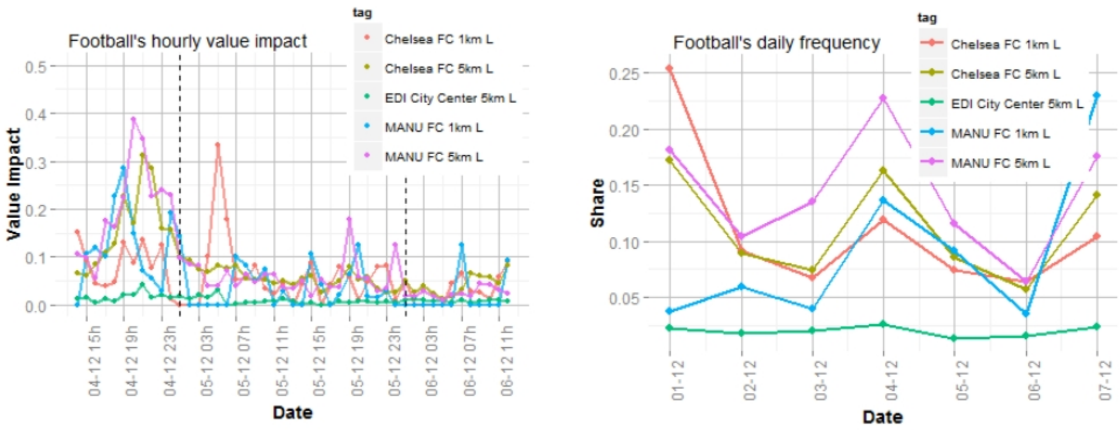

Figure 12: Football hourly detailed impact value (a) and frequency (b)

chosen the well-known RFM paradigm and introduced the concepts of exposure and engagement of a particular user with a topic to model the Monetary or Value component. Concretely in the industry domain, our new social media

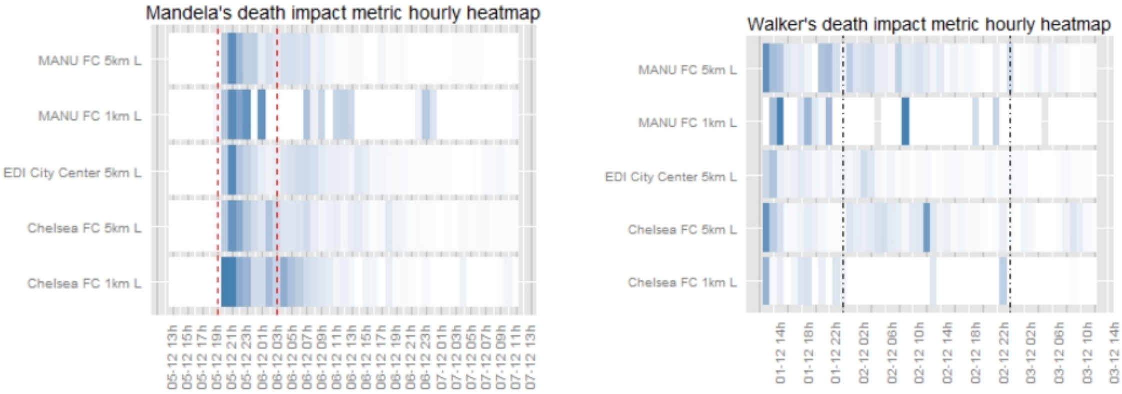

Figure 13: Mandela's (a) and Paul Walker's (b) death hourly impact heatmap

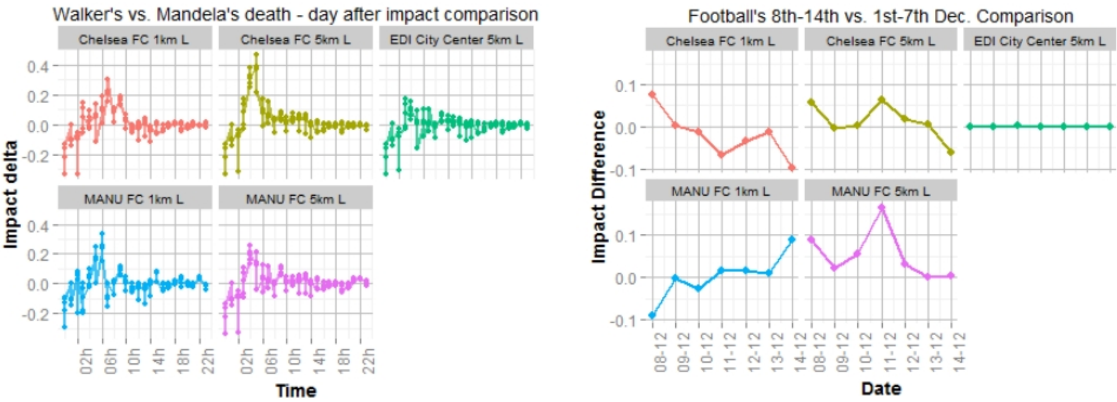

Figure 14: Mandela-Walker day after death comparison (a) and Week over Week football comparison (b)

RFM model could present a good performance in a variety of applications, ranging from event planning and marketing (campaign monitoring, topic affinity advertising, interest targeting) to market research (media monitoring, geo-located panelists, news impact).

The introduction of the exposure and engagement metrics allows for modelling at user level both passive and active topic impact and allows for filtering and segmentation based on different user attributes as all the metrics are defined at user level. As show-cased with the implemented system and the football analysis, our metrics perform well even in hourly chunks; they are consistent over time (delivering similar results in similar situations in different periods) and easy to understand (as they reflect the nature of the social network but at individual level).

In a variety of scenarios or extreme cases, the social media RFM model is proven to be robust always delivering meaningful metrics as discussed and demonstrated with the examples we analyzed based on the system that has also been implemented as part of this paper. One of the strengths of the approach we suggest in this work is the fact that the topic impact comparison is supported in heterogeneous scenarios, for example with different topics over different time frames in different locations.

For the sake of simplicity, in our social media model we have considered the links between users as equally powerful in the exposure calculation, which leaves the door open for improvement as different users might have differential influence power –popularity– on others which might result into a more realistic exposure result. Another challenging aspect inherent to the source of the data itself in the suggested approach is the bias derived from using online-only social network users to assess the topic impact on a particular place. The fact that we focus on English speaking users only reinforces the bias, as a topic might well impact different cultures and nationalities with different intensity. These limitations certainly point to future research directions.

Acknowledgements

This paper has been developed with the financing of projects TIN2010-17876, TIC5299, TIC-5991 and TIN2013-40658-P.

AlSumait, L., Barbará, D., Gentle, J., and Domeniconi, C. (2009). Topic significance ranking of lda generative models. In *Machine Learning and Knowledge Discovery in Databases*, pages 67–82. Springer.

Backstrom, L., Sun, E., and Marlow, C. (2010). Find me if you can: improving geographical prediction with social and spatial proximity. In *Proceedings of the 19th international conference on World wide web*, pages 61–70. ACM.

Barsade, S. G. (2002). The ripple effect: Emotional contagion and its influence on group behavior. *Administrative Science Quarterly*, 47(4):644–675.

Bennett, S. (2012). Twitter now has more than 140 million active users sending 340 million tweets every day.

Berners-Lee, T., Hendler, J., Lassila, O., et al. (2001). The semantic web. *Scientific american*, 284(5):28–37.

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022.

Bogart, L. (1967). *Strategy in advertising*. Harcourt, Brace & World New York.

Bult, J. R. and Wansbeek, T. (1995). Optimal selection for direct mail. *Marketing Science*, 14(4):378–394.

Cataldi, M., Di Caro, L., and Schifanella, C. (2010). Emerging topic detection on twitter based on temporal and social terms evaluation. In *Proceedings of the Tenth International Workshop on Multimedia Data Mining*, page 4. ACM.

Cavusoglu, H., Hu, N., Li, Y., and Ma, D. (2010). Information technology diffusion with influentials, imitators, and opponents. *J. Manage. Inf. Syst.*, 27(2):305–334.

Centola, D. (2010). The spread of behavior in an online social network experiment. *science*, 329(5996):1194–1197.

Cha, M., Haddadi, H., Benevenuto, F., and Gummadi, K. (2010). Measuring user influence in twitter: The million follower fallacy. In *4th International AAAI Conference on Weblogs and Social Media (ICWSM)*.

Cheng, Z., Caverlee, J., Lee, K., and Sui, D. Z. (2011). Exploring millions of footprints in location sharing services. *ICWSM*, 2011:81–88.

Chuang, J., Ramage, D., Manning, C., and Heer, J. (2012). Interpretation and trust: designing model-driven visualizations for text analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 443–452. ACM.

Clark, E., Roberts, T., and Araki, K. (2010). Towards a pre-processing system for casual english annotated with linguistic and cultural information. In *Proceedings of the Fifth IASTED International Conference*, volume 711, pages 044–84.

Deerwester, S. C., Dumais, S. T., Landauer, T. K., Furnas, G. W., and Harshman, R. A. (1990). Indexing by latent semantic analysis. *JASIS*, 41(6):391–407.

Fader, P. S., Hardie, B. G., and Lee, K. L. (2005). Rfm and clv: Using iso-value curves for customer base analysis. *Journal of Marketing Research*, pages 415–430.

Farris, P. W., Bendle, N. T., Pfeifer, P. E., and Reibstein, D. J. (2010). *Marketing metrics: The definitive guide to measuring marketing performance*. Pearson Education.

Fujita, S. and Fujino, A. (2013). Word sense disambiguation by combining labeled data expansion and semi-supervised learning method. *ACM Transactions on Asian Language Information Processing (TALIP)*, 12(2):7.

Griffiths, T. L., Steyvers, M., and Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological review*, 114(2):211.

Guille, A. and Hacid, H. (2012). A predictive model for the temporal dynamics of information diffusion in online social networks. In *Proceedings of the 21st International Conference Companion on World Wide Web*, pages 1145–1152. ACM.

Hughes, A. (2005). *Strategic Database Marketing*. New York: McGraw-Hill.

Jones, K. S. (1972). A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 28(1):11–21.

Khajvand, M., Zolfaghar, K., Ashoori, S., and Alizadeh, S. (2011). Estimating customer lifetime value based on rfm analysis of customer purchase behavior: Case study. *Procedia Computer Science*, 3(0):57 – 63.

Kiss, T. and Strunk, J. (2006). Unsupervised multilingual sentence boundary detection. *Comput. Linguist.*, 32(4):485–525.

Kumar, V. (2008). Customer lifetime value the path to profitability. *Foundations and Trends(R) in Marketing*, 2(1):1–96.

Kwak, H., Lee, C., Park, H., and Moon, S. (2010). What is twitter, a social network or a news media? In *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, pages 591–600, New York, NY, USA. ACM.

MacKay, D. J. and Peto, L. C. B. (1995). A hierarchical dirichlet language model. *Natural language engineering*, 1(3):289–308.

Manaris, B. (1998). Natural language processing: A human-computer interaction perspective.

Mathioudakis, M. and Koudas, N. (2010). Twittermonitor: Trend detection over the twitter stream. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*, SIGMOD '10, pages 1155–1158, New York, NY, USA. ACM.

Miller, G. A. (1995). Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.

Naaman, M., Becker, H., and Gravano, L. (2011). Hip and trendy: Characterizing emerging trends on twitter. *Journal of the American Society for Information Science and Technology*, 62(5):902–918.

O'Connor, B., Krieger, M., and Ahn, D. (2010). Tweetmotif: Exploratory search and topic summarization for twitter. In *ICWSM*.

Page, L., Brin, S., Motwani, R., and Winograd, T. (1998). The pagerank citation ranking: Bringing order to the web. In *Proceedings of the 7th International World Wide Web Conference*, pages 161–172, Brisbane, Australia.

Porcel, C., A., T.-L., Martnez, M., and Herrera-Viedma, E. (2012). A hybrid recommender system for the selective dissemination of research resources in a technology transfer office. *Information Science*, 184(1):1–19.

Rajyalakshmi, S., Bagchi, A., Das, S., and Tripathy, R. M. (2012). Topic diffusion and emergence of virality in social networks. *CoRR*, abs/1202.2215:1212–2232.

Romero, D. M., Galuba, W., Asur, S., and Huberman, B. A. (2011a). Influence and passivity in social media. In *Proceedings of the 20th International Conference Companion on World Wide Web*, WWW '11, pages 113–114, New York, NY, USA. ACM.

Romero, D. M., Meeder, B., and Kleinberg, J. (2011b). 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*, WWW '11, pages 695–704, New York, NY, USA. ACM.

Sakaki, T., Okazaki, M., and Matsuo, Y. (2010). Earthquake shakes twitter users: Real-time event detection by social sensors. In *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, pages 851–860, New York, NY, USA. ACM.

Salton, G. and McGill, M. J. (1986). *Introduction to Modern Information Retrieval*. McGraw-Hill, Inc., New York, NY, USA.

Scellato, S., Noulas, A., Lambiotte, R., and Mascolo, C. (2011). Socio-spatial properties of online location-based social networks. *ICWSM*, 11:329–336.

Sohrabi, B. and Khanlari, A. (2007). Customer lifetime value (clv) measurement based on rfm model. *Iranian Accounting & Auditing Review*, 14(47):7–20.

Stieglitz, S. and Dang-Xuan, L. (2013). Emotions and information diffusion in social mediasentiment of microblogs and sharing behavior. *Journal of Management Information Systems*, 29(4):217–248.

Sutherland, I. E., Sproull, R. F., and Schumacker, R. A. (1974). A characterization of ten hidden-surface algorithms. *ACM Comput. Surv.*, 6(1):1–55.

Tejeda-Lorente, A., Porcel, C., Peis, E., Sanz, R., and Herrera-Viedma, E. (2014). A quality based recommender system to disseminate information in a university digital library. *Information Science*, 261:52–69.

White, D. (2013). Social media growth 2006-2012.

Yang, Q. X., Yuan, S. S., Zhao, L., Chun, L., and Peng, S. (2003). Faster algorithm of string comparison. *Pattern Analysis & Applications*, 6(2):122–133.

Ye, S. and Wu, S. F. (2010). Measuring message propagation and social influence on twitter.com. In *Proceedings of the Second International Conference on Social Informatics*, SocInfo'10, pages 216–231, Berlin, Heidelberg. Springer-Verlag.

1.2 Quantifying the emotional impact of events on locations with Social Media

- Bernabé-Moreno, J., A. Tejeda-Lorente, C. Porcel, Fujita.H and E. Herrera-Viedma. "Quantifying the emotional impact of events on locations with Social Media" Knowledge Based System.
  - Status: **Submitted for publication**.
  - Impact Factor (JCR 2013): 1.977
  - Subject Category: Computer Science, Artificial Intelligence. Ranking 17 / 688 (**Q1**).
  - Subject Category: Computer Science, Information Systems and Management. Ranking 10 / 122 (**Q1**).
  - Subject Category: Computer Science, Management Information Systems. Ranking 6 / 76 (**Q1**).
  - Subject Category: Computer Science, Software. Ranking 32 / 1193 (**Q1**).

Quantifying the emotional impact of events on locations with Social Media

J. Bernabé-Moreno[a,*], A. Tejeda-Lorente[a], C. Porcel[b,*], H. Fujita[c], E. Herrera-Viedma[a,*]

[a] *Department of Computer Science and A.I., University of Granada, Granada E-18071, Spain*
[b] *Department of Computer Science, University of Jaén, Jaén E-23071, Spain*
[c] *Iwate Prefectural University (IPU), Iwate, Japan*

Abstract

Today's digital world is difficult to conceive without the existence of Social Media. An ever growing part of our daily communicative activity takes place in the Social Media platforms, where not only *what* we say is kept, but also *when* and increasingly *where* we say it. The *how* we communicate is also very insightful, as the words we chose in our communication reveal a lot about our emotions. Inspired by these ideas, we created a new method to quantify the emotional impact of an event on a particular location in absolute terms but also broken down to the different emotional states. To support that, we explored different modelling approaches for the emotional profiling of locations adopting the well established Pleasantness-Arousal-Dominance paradigm. In this paper, we define our method, explain how we extract the emotions from Social Media Interactions relying on a modified version of extended Affective Norms for English Words, describe the system we implemented for the method's validation and discuss the overall performance of our approach with different emotionally rich events in three known locations.

Keywords:
Emotional Impact; Emotional Profiling; Social Media;Emotional Mapping

1. Introduction

Social Media (SM) is increasingly becoming an important part of our lives in a more and more integrative way. If we have a look at the vision statement of Twitter *"Our mission: To give everyone the power to create and share ideas and information instantly, without barriers"* [1] we realize that this vision is actually not very far from the reality any more.

As internet became pervasive with the advent of mobile and wireless technologies -such as Universal Mobile Telecommunications System (UMTS), Long-Term Evolution (LTE) and Wireless Fidelity (WiFi)-, posting SM updates or consuming SM content was no longer limited to those sitting in front of a PC with wired access to the World Wide Web. Mobile connectivity took SM to a whole new level and brought Twitter's vision one step closer to its realization by making the "instant" component actually feasible. As a consequence of that, the location where the interactions took place increasingly became an integral part of the SM dialogue. The geo-tagging of the SM interactions started to be supported by the traditional SM platforms and new platforms emerged, where the role of the location surpassed the content itself, such as Foursquare [2], that provides personalised local search experience for its users by taking into account the places a user goes, the things they have told the app that they like, and the other users whose advice they trust. As a result, the proportion of SM interactions that in addition to the known *time-stamp* got a *location-stamp* started to increase drastically, opening the door to a whole new set of insights for a location analysis based up the SM users and the SM interactions tagged in the location [1, 2]. The accuracy of the geo-location tags

*Corresponding author. Tel.: +34-958-24-4258 ; fax: +34-958-24-3317 .
E-mail address: juan.bernabe-moreno@webentity.de,cporcel@ujaen.es,viedma@decsai.ugr.es.

[1] `https://about.twitter.com/company`
[2] `https://foursquare.com`

could vary from a few meters in the case of GPS powered pair of latitude-longitude geographical coordinates to the name of a district, a known place or even a city, supporting different kinds of analysis.

One of the key success factors of the rapid SM adoption is the democratization of the digital media: with initiatives such as the blogosphere [3], everybody could make their own contributions to the content published by any author, anybody could become an author and engage with others in a digital dialogue or anybody could find, read and participate in any existing SM conversation. The SM platforms based on the concept of micro-blogging took it to the next level, as everybody could be an author and a reader any time. The push-first, comment-later paradigm so popular in the blogosphere started to look old-fashioned. Rather, anybody was empowered to initiate a communication, enrich an existing thread, jump from a thread to another one, ignore, criticize, share richer content like pictures, videos, etc. The ease of publishing, sharing and consuming content boosted the adoption of these platforms as the place to talk any time about everything with everybody. Users also became less reluctant to express -almost in a unfiltered way- what's literally going through their minds [4] in micro-blogging sites, unlike other purpose-specific SM platforms –such as LinkedIn, etc–. As a side effect, the amount of information generated in SM drastically increased introducing the need for recommendation systems to separate the relevant content from the rest [5, 6].

The content generated in the SM interactions has been subject of prolific research in the recent years. Sophisticated machine learning methods to estimate or extract emotions from the content created by users has been developed [7], including support vector machines [8], bayesian networks [9], maximum entropy approaches [10] and concept-level analysis of natural language text [11] supported by combinations of common-sense-reasoning [12] and ways of representing affection, such as affective ontologies [13], etc-. The approaches mentioned above require longer high-quality text to work properly. These criteria cannot be met by the kind of interactions created in micro-blogging platforms, because of following reasons: posts are typically short –e.g.: Twitter doesn't allow for posts longer than 140 characters–, disconnected from each other –appending subsequent posts is rarely a viable option– and with a lot of abbreviations, spelling mistakes, etc. To tackle this problem, further propositions based on affective dictionaries where the emotional rating of each word could be looked up were explored. One of the most popular approaches is the Affective Norms for English Words (ANEW) [14], consisting of a pre-defined set of 1034 frequently used English words that have been rated using the so called Self-Assessment Manikin [15]. A randomly selected group of people were asked to read a corpus and to provide the rating for each occurrence of these words. The resulting dictionary contained three statistically normalized -mean and standard deviation- scores for each word corresponding to the three PAD emotional state model components created by Mehrabian et al. [16] back in 1980. These components are *(P)leasure* or valence –the pleasantness of the emotion–, *(A)rousal* –the intensity of emotion provoked by the stimulus– and *(D)ominance* –the degree of control exerted by the emotion–. For example, *fear, rage, anger* and *boredom* are all unpleasant emotions, but *rage* is clearly more aroused and more intense than *boredom* and *fear* is rather submissive in contrast to *anger*. With complex emotion representation models like the one suggested by J. Russell in 1980 [17], the $valence, arousal$ pair could be mapped to particular named emotions or moods.

In this paper, we want to explore the potential of these ideas combined, namely the emotions extraction from SM posts and the ability to geo-locate SM interactions to provide unprecedented insights for locations. Thus, the purpose of this piece of work is defining a method to quantify the emotional impact of different events during a period of time on a given location based on the SM user generated content attached to this location. The method we are suggesting here pursues the creation of emotional profiles for locations by extracting and normalizing the emotional payload of the SM interactions created in the location over time. These profiles serve as reference or "norm" to assess the emotional pattern of a particular event against. In [18], the authors came up with a compelling analogy to consider that personality is to emotion as climate is to weather: what one expects is personality, what one observes at any particular moment is emotion. Our method measures the divergence between the emotional profile of what's happening during a given period of time –the weather– from the emotional baseline or emotional profile –the climate– of the location.

Emotional profiling of locations in general and emotional impact measuring in particular open a new door to marketing activities. Choosing the right marketing message that fits best the emotional baseline of a location can drastically impact the performance of a campaign. Understanding the local impact of different event types makes the identification of promotional activities easier [19]. Political campaigns could also rely on this kind of insights to chose the right wording in their massages and then measure the outcome using our approach even before the elections have taken place. At a particular level, a person might be also interested in understanding how good his/her personality matches the emotional profile of a potential place to move to. These are just a few examples of the countless applications of the output of this paper.

This paper is organized as follows: firstly we provide all the background information relevant for our research. Then we introduce our method to create emotional profiles of locations and to quantify the emotional impact of a particular event on the location comparing the location profile with the event's emotional footprint. After that, we extensively describe the system we implemented to demonstrate our method with real-world data and subsequently we show some practical examples to discuss the performance of our emotional impact quantification method. We finalize our paper sharing our conclusions and pointing out future research lines to take forward this piece of work.

2. Background and related work

Emotional models and affective architectures have been intensively researched in the last 15 years in all variety of fields, such as Artificial Intelligence, Human-Computer Interaction, Robotics, Gaming, etc [20]. Yet the first attempts to create a model to compare emotional states were made in the cognitive sciences domain. At an early stage of development, the intensity –or arousal–, the degree of pleasantness –valence– and the amount of influence you feel the environment has upon you –dominance–, were explored independently and represented with different scales. Based on the work initiated in [21] and [16] where the *Pleasure-Arousal-Dominance* (PAD) model was formally introduced, Russell suggested in a seminal work the combination of emotional axis to create a circumplex model that enabled the position of emotions on a plane [17]. For the representation of each emotional state, Russell suggested a pair of coordinates on a two dimensional space: on the x-axis the valence and on the y-axis the arousal of the stimulus. Up to 28 emotional states have been multidimensionally scaled in Russell's model, so that intermediate terms are polar opposites (e.g.: excited-depressed, distressed-relaxed, etc). Several new models and refinements on Russell's model followed, each one conceptualizing the dimensions in different ways: tension and energy [22], positive and negative affect [23], approach and withdrawal [24], etc.

Bradley and Lang created in 1999 [14] a set of normative emotional ratings for 1034 commonly used English words, also known as the set of Affective Norms for English Words or ANEW. Based on the outcome of this research, it was possible for the first time measuring natural language fragments in terms of the PAD model dimensions. This seminal work can be considered the first enabler for the emotional states extraction from user generated content. Fourteen years later, an extended version of ANEW (eANEW) containing more than 13K English lexemes and faceted by gender and education level was developed applying almost the same procedure as in the original piece of work [25]. ANEW became very popular in the research community and was adapted for other languages [26, 27, 28, 29]. In addition to ANEW, further affective dictionaries have been created, for example Affective Wordnet [30], where semantic synsets are assigned one or several affective labels for those concepts representing moods, situations eliciting emotions, or emotional responses.

The recent years have witnessed the creation of countless approaches to extracting emotional states from user generated content. In [31] Ramaswamy et al. created an interactive tool to visualize the emotions extracted from a Twitter query over the Russell's 2D plane for the most recent time. This tool also allows for a keyword extraction based on frequency, as well as the visualization of a moods' heatmap over time. In [32], the authors went even further and mapped the emotions to a 3D virtual human. An additional interesting contribution of this paper is the color interpretation of the emotions mapping different values of arousal and valence to colors. In [33] the authors analysed the role of the different emotional states in the information diffusion in SM. In [34] the authors explored the emotions distribution over 2.5 million post in the BBC forum, analysing the correlation between negative emotions and users activity. In [35] a predictive model for blog posts ratings providing the estimated level of valence and arousal of a post on a ordinal scale was presented, also taking as a basis the Russell's circumplex model. In [36] the authors created near real-time, remote-sensing, non-invasive, hedonometer consuming geo-localized tweets from the Twitter Streaming API. The happiness extraction from the micro-posts relies on an own crowd-sourcing effort where over 10000 words were rated for happiness, instead of adopting the traditional ANEW family.

The analysis of geo-localized SM user generated content to augment the knowledge and gain new insights about a location has also been object of intense research. In [37] the authors present a framework and set of metrics to quantify the impact of a topic in different locations adapting the Recency-Frequency-Monetary paradigm. The work created in [19] presents an innovative approach to extract near-real time information from SM related to a brand to enrich customer acquisition and retention campaigns. In [38], the authors suggest the real-time analysis of localized interaction to implement an early warning system. In the area of disasters prevention, geo-localized interactions have been intensively analysed to create early warning and prediction systems on natural catastrophes such as earthquakes, tsunamis, etc. in particular locations [39, 40, 41]

3. Defining a new method for quantifying emotional impact

In this section we proceed with the formal definition of our approach for measuring the emotional impact of an event on a particular location.

3.1. Preliminary definitions

To support the metrics definition in our methodology, we first introduce a set of relevant concepts:

Definition 1. *The set U represents the set of social media users from which we have evidence they have been in the location L we are monitoring during the time period under analysis* $\triangle t$

$$U \equiv \{u\},\ \forall u_i \in U,\ InLocation(u_i, L, \triangle t) \tag{1}$$

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

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

$Follows(u_i, u_j)$ 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 symmetric; although in several SM platforms it is the case (e.g.: Facebook or Linked.in).

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

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

Definition 4. *We define all user interactions* $(Interactions)$ *for a given user* u_i *over a time interval* $\triangle t$*, as:*

$$Interactions\ (u_i, \triangle t) \equiv \{it\},\ \forall it_i \in Interactions\ (u_i, \triangle t),\ Author(u_i, it_i, \triangle t) \tag{4}$$

A Social Media Interaction represents the atomic piece of content generated by the user u_i during the time $\triangle t$ in a Social Media Platform (e.g.: a tweet, a re-tweet). Thus, $Author(u_i, it_i, \triangle t)$ is a function that retrieves $True$ if u_i created the interaction it_i in the time period $\triangle t$, 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*.

An SM interaction it_i can be also seen from the Natural Language perspective as a set of terms $terms(it_i)$: $it_i \equiv \{t\}, \forall t_j, t_j \in T$ where T represents all possible terms in the English language, including spelling mistakes, newly invented terms and whatever communication unit which conveys a meaning between the sender and at least one of the recipients.

3.2. Modelling emotions

We provide now a set of definitions to formally describe the emotional state Pleasantness-Arousal-Dominance model in the SM context.

Definition 5. *We define the emotional rating ER of a user interaction* it_i *as a vector with three components: valence v, arousal a and dominance d*

$$ER(it_i) \equiv [v, a, d] \tag{5}$$

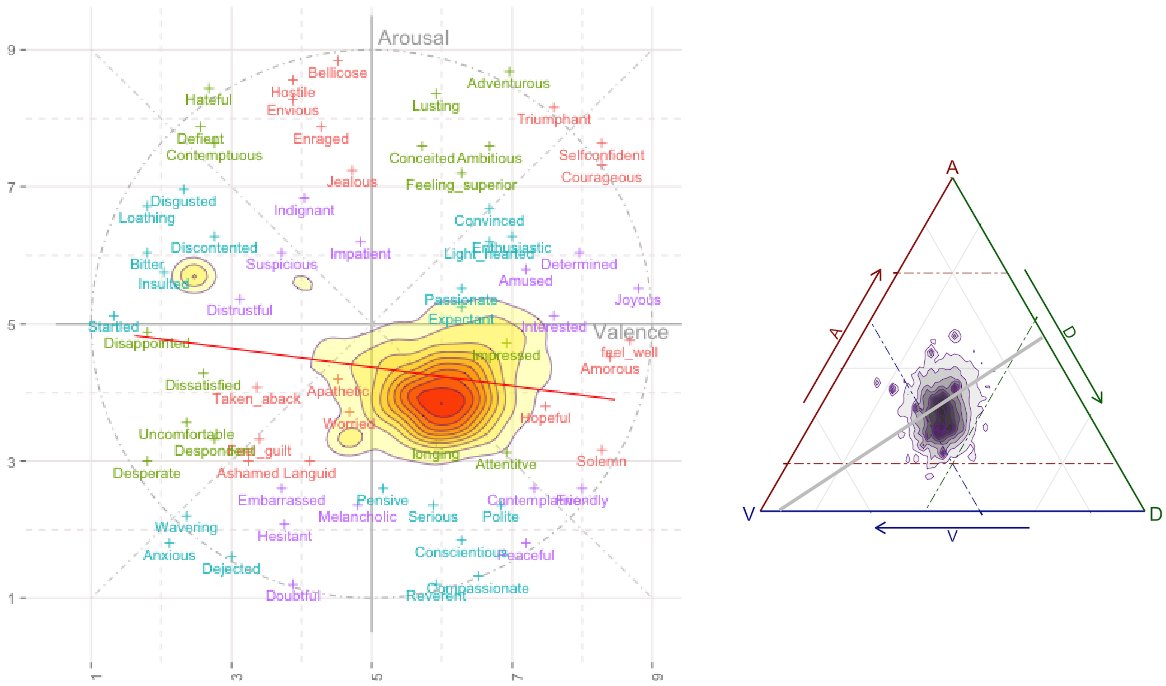

Figure 1: (a) Two-dimensional circumplex space model created in [17] and refined in [42, 32, 35] with a sample baseline distribution (b) Three dimensional Valence-Arousal-Dominance model with a sample baseline distribution

To obtain the values for valence v or pleasantness, arousal a and dominance d, our approach relies on a set of aggregated rating functions defined on top of the extended version of the ANEW lemmatization (eANEX) ([25]). Each PAD component in the vector is obtained applying a function that looks up the interaction lexemmas in the eANEW dictionary, retrieves the rating values for each available one and combines the results into a single value with a weighted average operation. As the eANEW also provides for each rated lexemma the standard deviation for all the raters, we use the maximum probability value assuming a normal distribution as the weight for each lexemma $f_{max} = \frac{1}{\sigma\sqrt{2\pi}}$ to give higher weight to rating with lower sparsity.

Thus, a generic rating function is defined as follows:

$$r(it_i) = \frac{1}{\sum_{j=1}^{|terms(it_i)|} f_{max}(t_j)} \sum_{j=1}^{|terms(it_i)|} \rho(t_j) * f_{max}(t_j),\ t_j \in terms(it_i) \tag{6}$$

where $\rho(t_j)$ can be the eANEW valence mapping $\upsilon(t_j)$ to obtain v, or the eANEW arousal mapping $\alpha(t_j)$ to obtain a or the eANEW dominance mapping $\delta(t_j)$ to obtain d.

To translate the values of v, a, d to named emotional states, we make use of the enhanced adaption of Russell's circumplex model as showed in Fig. 1 (a), which only rely on 2 components, valence and arousal. Motivated by the defence of the usefulness of dominance measuring emotions in [43], we also explore in our Emotional Impact calculation a model with all three components (see Fig. 1 (b)).

Definition 6. *We define the emotional baseline EB of a location L over a given period of time $\triangle t$ as a valence-arousal-dominance distribution resulting from the aggregation of all interactions' emotional ratings authored by the users in the location during the period of time $\triangle t$*

$$EB(L, \triangle t) \equiv [v, a, d] \tag{7}$$

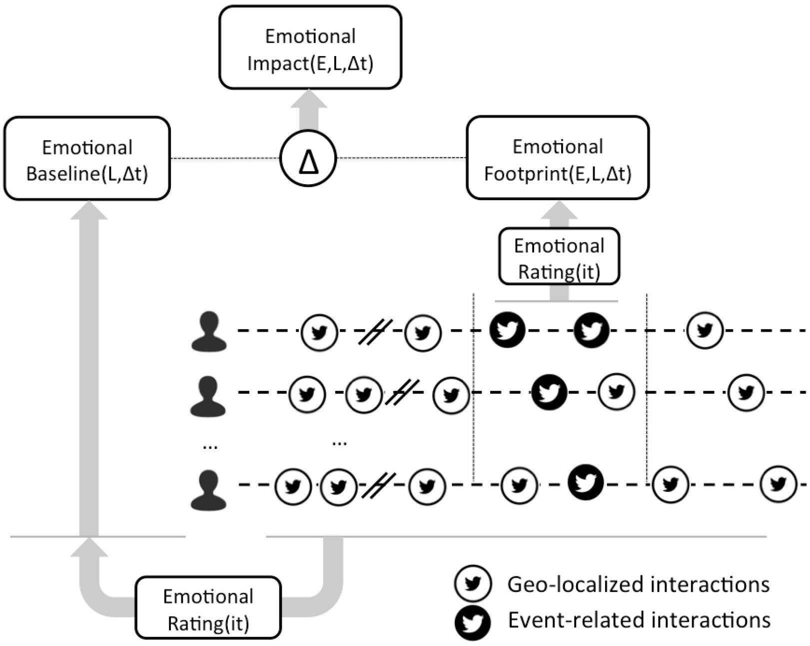

Figure 2: Emotional Impact metrics orchestration

where $[v, a, d] = \Im(ER(it_j))$, $Author(u_j, it_i, \triangle t)$, $u_j \in U$, $InLocation(u_j, L, \triangle t)$ The function $\Im$ can be designed to give more weight to interactions more recent in time, for example to adjust to potential personality changes in individuals.

To model the distribution of emotions in the emotional plane, we suggest a multivariate kernel density function ([44]), defined as follows:

$$\iota_H(x) = \frac{1}{n} \sum_{i=1}^{n} K_H(x - x_i) \tag{8}$$

where $x = (x_1, x_2, , x_d)^T, x_i = (x_{i1}, x_{i2}, , x_{id})^T, i = 1, 2, , n$ are the ER vectors; H is the bandwidth (or smoothing) matrix (chosen as described in [45]); K is the kernel function which is a symmetric multivariate density; $K_H(x) = |H|^{1/2} K(H^{1/2} x)$

An additional implementation for emotional baseline $EB(L, \triangle t)$ could be coupled to time chunks to incorporate the seasonality effects. The time granularity level depends on the variability for the particular location. Thus, one could create a baseline for a given *month of the year* -e.g.: December because of Christmas is different than February in places where Christmas is important... or the Ramadan month vs. an ordinary one in Islamic countries, etc-, *day of the week* -e.g.: a Monday vs. a Friday- or even *hour of the day* -eg.: 10:00h vs. lunch time-.

3.3. Emotional Impact

Once we have an emotional baseline for a location, we can define the metrics for assessing the impact of an event on a location as a deviation from the baseline.

In order to do that, we need to obtain the *emotional footprint* of the event in the location, which follows the same procedure as we defined to obtain the baseline for the place, just for the subset of interactions related with the event.

Definition 7. *We define the set of Interactions for a given user u_i with the event E over a time interval $\triangle t$ as:*

$$Interactions\ (u_i, E, \triangle t) \equiv \{it\},\ \forall it_i \in Interactions\ (u_i, \triangle t),\ Author(u_i, it_i, \triangle t)\ \wedge related(it_i, E) \tag{9}$$

Where $related(it_i, T)$ is a NLP membership function retrieving $True$ if the iteration it_i is connected to the topic T –intuitively, one or more words from the semantic field for the topic T are mentioned in it_i– and $False$ otherwise.

Definition 8. *We call emotional footprint EF of a given event, to the aggregation of the emotional ratings of the interactions related to this event over a period of time $\triangle t$*

$$EF(L, E, \triangle t) \equiv [v, a, d] \tag{10}$$

Typically, the aggregation function is the same $\Im(ER(it_j))$ we used for the emotional base-lining of the location L in Def. 6.

Once we have all ingredients in place, we can define the emotional impact of an event on a place as the difference between the location's emotional baseline $[v, a, d]$ distribution and the event's emotional footprint $[v, a, d]$ distribution, as explained in Fig. 2:

$$EI(L, E, \triangle t) \equiv \frac{|It(L, E, \triangle t)|}{|It(L, \triangle t)|} |EB(L, \triangle t) - EF(L, E, \triangle t)| \tag{11}$$

As we employed multivariate kernel density functions for modelling both $EB(L, \triangle t)$ and $EF(L, E, \triangle t)$, to quantify the difference we suggest applying the standard deviation of the resulting difference distribution: $\sigma(|EB(L, \triangle t) - EF(L, E, \triangle t)|)$. $\frac{|It(L,E,\triangle t)|}{|It(L,\triangle t)|}$ represents the share of interactions related to the event vs. the whole set of interactions that have been gathered and thus making the impact dependant on the portion of activity related to the event.

Mood	Valence	Arousal	Mood	Valence	Arousal	Mood	Valence	Arousal	Mood	Valence	Arousal
Triumphant	7.60	8.16	Amused	7.20	5.80	Startled	1.32	5.12	Despondent	2.76	3.32
Selfconfident	8.28	7.64	Joyous	8.80	5.52	feel_well	8.68	4.76	Desperate	1.80	3.00
Courageous	8.28	7.32	Interested	7.60	5.12	Amorous	8.40	4.52	Friendly	8.00	2.60
Adventurous	6.96	8.68	Convinced	6.68	6.68	Hopeful	7.48	3.80	Contemplative	7.32	2.60
Lusting	5.92	8.36	Light_hearted	6.68	6.20	Solemn	8.28	3.16	Peaceful	7.20	1.80
Conceited	5.72	7.60	Enthusiastic	7.00	6.28	Impressed	6.92	4.72	Polite	6.84	2.36
Feeling_superior	6.28	7.20	Passionate	6.28	5.52	longing	5.92	3.28	Conscientious	6.28	1.84
Ambitious	6.68	7.60	Expectant	6.28	5.24	Attentive	6.92	3.12	Compassionate	6.52	1.32
Bellicose	4.52	8.84	Indignant	4.04	6.84	Apathetic	4.52	4.20	Reverent	5.92	1.20
Hostile	3.88	8.56	Impatient	4.84	6.20	Worried	4.68	3.72	Serious	5.88	2.36
Envious	3.88	8.28	Suspicious	3.72	6.04	Feel_guilt	3.40	3.32	Pensive	5.16	2.60
Enraged	4.28	7.88	Distrustful	3.12	5.36	Languid	4.12	3.00	Melancholic	4.80	2.36
Jealous	4.72	7.24	Disgusted	2.32	6.96	Ashamed	3.24	3.00	Embarrassed	3.72	2.60
Hateful	2.68	8.44	Loathing	1.80	6.72	Taken_aback	3.36	4.08	Hesitant	3.76	2.08
Defiant	2.56	7.88	Discontented	2.76	6.28	Disappointed	1.80	4.88	Doubtful	3.88	1.20
Contemptuous	2.76	7.64	Bitter	1.80	6.04	Dissatisfied	2.60	4.28	Wavering	2.36	2.20
Determined	7.96	6.04	Insulted	2.04	5.76	Uncomfortable	2.36	3.56	Anxious	2.12	1.80

Table 1: Circumplex Moods Mapping

We can enhance this overall impact quantification by defining an impact metric at named mood level. For that, we need to provide an additional definition on top the emotional rating of an interaction to assign a named mood.

Definition 9. *We define the set of named moods NM as the set of emotional states available in the extended Circumplex Model, each one with a pair of "valence, arousal" coordinates.*

The Circumplex Model was first created in [17] and refined and extended in [32, 35, 42]. The set of named moods as well as their $[v, a]$ coordinates can be seen in the Table 1.

Definition 10. *We define the leading mood of a user interaction it_i to the closest named mood to the valence v and arousal a components of the emotional rating of the $ER(it_i)$*

$$LeadingMood(it_i) \equiv m_j, m_j \in NM, m_j = \min_{\forall m_k \in NM} dist_{[v,a]}(it_i, m_k) \tag{12}$$

Based on this definition, both Location Emotional Baseline and Event Emotional Footprint can be expressed as the share of each named mood being *Leading Mood* during the period of time under analysis. For example, if we had one event with 40 interactions with following leading moods: 20 *longing*, 10 *pensive* and 10 *interested*, the share would be 0.5 *longing*, 0.25 *pensive* and *interested*.

It allows us to define a new version of the emotional baseline metric for a Location in terms of a particular named mood as follows:

$$EB_{NM}(m_j, L, \triangle t) \equiv \frac{|LeadingMood(It(L, \triangle t)) \cap \{m_j\}|}{|It(L, \triangle t)|} \tag{13}$$

The emotional footprint of an event referred to a particular named mood can also be defined in a similar way:

$$EF_{NM}(m_j, L, E, \triangle t) \equiv \frac{|LeadingMood(It(L, E, \triangle t)) \cap \{m_j\}|}{|It(L, E, \triangle t)|} \tag{14}$$

Based on these new metrics, we then provide a named mood version of the impact quantification as follows:

$$EI_{NM}(m_j, L, E, \triangle t) \equiv EF_{NM}(m_j, L, E, \triangle t) - EB_{NM}(m_j, L, \triangle t) \tag{15}$$

Intuitively, this metric represents how a particular mood become more or less important –share increase or decrease– in the event emotional footprint versus the location emotional norm.

In the subsequent sections we are going to provide a description of the system we propose to implement these metrics and discuss their performance with the help of a real-world example. The reader is going to get more clarity about the definition and the usage of the set of equations we just presented.

4. System Architecture

The purpose of this section is describing the system we have built to implement the metrics defined in our method for the emotional impact quantification of events on locations.

The set of metrics we just defined are in principle platform agnostic. We've chosen Twitter to implement our system because of following reasons:

- Ease of information extraction: almost no restrictions to get a significant sample of all interactions providing a set of query parameters
- Text-based content dominance: unlike other platforms favouring more rich media content –videos, pictures, etc–.
- High share of geo-located interactions.
- High-engagement general purpose platform

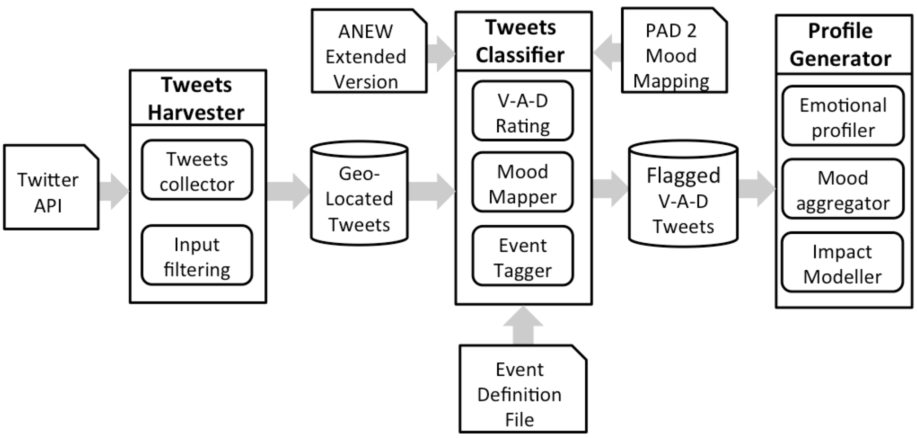

Figure 3: System architecture overview

The system technical architecture is based on the footprint explained in [46]. From the functional perspective, our system connects to the publicly available Twitter Search API [3] to poll the geo-located tweets for the location, consults the event definition file to flag the tweets related to the event, applies the eANEW emotional rating of the content and builds the emotional profile for the location and the emotional footprint for the event to finally produce the impact metrics described in the previous section.

The system consists of 3 different modules in charge of different labours all along the process. Each module is implemented by a set of components with a clearly defined function (see Fig. 3). In the following sections we are going to describe how the different modules work and what the role of the components being involved is.

4.1. Tweets harvester

The harvester collects all tweets created in a given area. An area is defined in our systems as a pair of geographical coordinates –latitude - longitude– and a radius. This module also applies a language filter to avoid the later emotional rating of non-English tweets, as we are working with the eANEW. In principle, the system could also work with Affective Norms for other languages which would also adjust the language filter of the harvester.

4.2. Tweets classifier

The purpose of this module is the flagging of the tweets related to the event, the emotional rating of the harvested tweets as well as the mood flagging, which is carried out by three components:

The *event flagger* marks all tweets related to the events The event definition file usually contains three types of information:

1. *Social Media Entities related to the topic*: Set of official accounts, nicknames, hashed tags, etc. users mention in their interactions with the event (e.g.: for a Roland Garros final tennis final match, we would have *RafaelNadal* for *Rafa Nadal*, *DjokerNole* for *Novak Djokovich*, etc). For completeness it should include both official accounts and those that are not official but with high levels of activity.
2. *Topic Named Entities*: set of named entities related to the topic (e.g.: *Rafael Nadal*, *Noval Djokovic*, etc)
3. *Topic Lexicon File*: containing the set of non-named entities related to the topic (e.g. in the tennis domain: *ace*, *match ball*, *set*, *advantage*, etc.)

Each geo-located tweet is tokenized applying a sentence tokenizer first and a word tokenizer later (based on [47]) both adapting the Punkt Tokenizer [48] to deal with social media texts. The modified tokenizer provides the stop words removal as well. The *event flagger* intents to match each and every reference term listed in the SM and Named Entities files applying a string similarity algorithm [49], which delivers a similarity score. The matching procedure implements thresholds -that may differ depending on the source– to support the fact that the social media content is often full of spelling errors [50], which is likely to happen even more frequently when it comes to named entities of foreign people (e.g. staying in tennis, *Nalbandian* is often spelled as *Nabandian* even by renowned tennis twitter accounts).

The *V-A-D rating component* lemmatizes the content of each tweet, performs the eANEW lookup and applies the weighting averaged defined in the Equation 6, providing a value for the valence, arousal and dominance. Some constraints can be applied to avoid volatile results when for example just one lemma out of the entire tweet content is found in the eANEW file. In this case, the system would produce an NA. Prior to the lemmatizion we apply a set of NLP components such as a sentence tokenizer followed by a word tokenizer (based on [47]) both adapting the Punkt Tokenizer [48] and a stemming algorithm to remove stop words, similar to the *event flagger*.

The *Mood mapper* assigns an emotional state to the resulting $[v, a]$ pair, applying a pre-defined moods mapping file (see [35]). Basically, it applies a refinement of the Russell's circumplex emotions model, as explained in Fig. 1. Each interaction represented by a pair of $[v, a]$ values is assigned to the Mood label –what we defined as Leading Mood in the Equation 12– whose circumplex coordinates are the closest to these $[v, a]$.

The result of applying the Tweets classifier is a set of tweets, each one with a $[v, a]$ score and a *mood* assigned.

[3] Available at `https://dev.twitter.com/docs/api/1/get/search`

4.3. *Profile Generator*

After the emotional rating, event flagging and moods mapping of the harvested SM interactions, this module aggregates the results into a location emotional profile on one hand and creates the emotional footprint for the event on the other hand.

The *Emotional Profiler* extracts a kernel density function (see Equation 8) in a bi-dimensional and tri-dimensional spaces with all $[v, a]$ and $[v, a, d]$ ratings respectively obtained from the previous steps. This function represents the emotional baseline profile of the Location L, as explained in the Subsection 3.2. The same procedure is applied to extract the 2D and 3D kernel density functions that represent the emotional event footprint with the subset of tweets flagged as related to the event.

The *Mood Aggregator* provides an aggregated view of the flagged moods collected over the time period in terms of absolutes and share for both the emotional baseline of the location and the emotional footprint of the event, as explained in the Subsection 3.3. This component produces the named mood versions of the location baseline and event footprint (see Equations 13 and 14).

The *Impact Modeller* quantifies the impact applying the Equation 11 as explained in 3.3 and providing also a quantification at named mood level.

The system we just described can be easily adapted to work with other languages. It would require adjusting the *Input filtering*component in the *Harvester* and replacing the Affective Norms definition file for English by the one of the target language in the *V-A-D Rating* component in the *Profile Generator*.

5. Evaluating our approach to quantify the emotional impact on locations

In this section we are going to show how three different real-world locations have been impacted by 2 tragic events that happened 6 days apart from each other and shook the hearts of multitudes within the space of one week. We are talking about the death of the famous American actor Paul Walker 5 on November the 30th 2013 and the decease of the charismatic Peace Nobel Price winner, South Africa's first black president and anti-apartheid icon Nelson Mandela 6 days later. We deliberately chose two events with the same tragic background to show the full potential of our method and demonstrate how different emotions can surface and how we are able to detect them.

5.1. *The set-up*

We set up 3 harvesters located in emblematic places in Great Britain cities: Manchester, centred on the Old Trafford Stadium (53.463101, -2.291490), in the popular Chelsea borough in London, centred on the FC Chelsea FC Stadium (51.481543, -0.190866) and in the Edinburgh City Center (55.9537,-3.188980), all three with a radius of 5 km. Thus, we covered a rather peripheral area of Manchester, and two pretty central areas of London and the Scottish capital... so quite different from each others.

The Event definition file for both events have been created with all named entities of both personalities and the popular aliases people use to refer to them (e.g.: *Madiba* for Nelson Mandela), to their contribution (e.g.: #2F2F hashtags for Walker's master piece *Too Fast, too Furious*) and combinations making reference to the sad incident (e.g.: *RIPPaul*).

Our harvesters ran for longer than 3 months, but we are going to focus our analysis on the first two weeks of December 2013, when both events manifested. The harvesters gathered 1088627 tweets during these 2 weeks in the mentioned locations. Applying a language filter (just "English") and the quality filter (just tweets with at least 2 words with eANEW rating), we ended up having 6522 tweets related to Walker's death and 6324 tweets to Mandela's death.

5.2. *The emotional impact quantification*

As explained all along this paper, the pre-requisite for the emotional impact quantification is the emotional base-lining of the locations and the creation of the event emotional footprint.

We have obtained them in two time-granularity levels: hourly and daily. Providing a hourly view over time helps us understanding the carousel of emotions that such a tragic event like the death of these two beloved personalities can trigger. In Fig. 4 we represent the event footprint (yellow–red gradient) vs. the location baseline (gray gradient) in the emotional circumplex plane of all three locations for a particular time, 10 o'clock of the day after the tragic incidents. In general, we appreciate a shift to the left –the "sad" quadrants–, with the forming of high-density centroids around different named moods depending on the location and the event:

- People in Chelsea are clearly *taken aback* by Mandela's death while talk *passionately* about it, expressing some *distrust* and some *dissatisfaction*.
- In Edinburgh, the reaction to Walkers death in comparison to Mandelas manifests in a more intense manner. Masses talk *passionately* to express their *dissatisfaction* and *discomfort*, with some doses of *bitterness* and *discontent*.
- Manchester follows the same pattern: Walkers death has a greater impact at this particular time, covering almost all negative moods in the mid-positive to mid-negative arousal spectrum (*dissatisfaction*, *discomfort*, *shame*, *discontent*, distrust, *bitterness*, etc).

If we incorporate to our analysis the Dominance component (see Fig. 5), we also observe a shifting triggered by both events in all locations, but following the same pattern as just discussed in the *valence*, *arousal* circumplex plane.

The emotions are very changing, that's why an impact quantification makes more sense at daily level; having more interactions related to the event –1 day vs. just 1 hour– makes the analysis results less volatile on one hand and changing emotions get to equalize along the day on the other hand. Nonetheless, the hourly change of the emotional footprint in both cases is of great interest for appreciating the variety of emotions that a tragic incident can release. Therefore, we have created 4 animations (bi-dimensional and tri-dimensional) where we show it hourly for the first days after both deaths and made them available in the popular SM platform YouTube.com (see Table 2)

Nelson Mandela's Death hourly emotional footprint vs. emotional baseline of our three locations under analysis.	`https://youtu.be/utqckiiYdUo`	
Paul Walker's Death hourly emotional footprint vs. emotional baseline of our three locations under analysis.	`https://youtu.be/jw0eZbPRki0`	
Hourly Emotional Location baseline vs. Mandela's death emotional footprint in a 3D PAD plane	`https://youtu.be/kTKxoAb65no`	
Hourly Emotional Location baseline vs. Walker's death emotional footprint in a 3D PAD plane	`https://youtu.be/eX-F-g_v9yg`	

Table 2: Animated emotional location base-lining and event foot-printing

In Fig. 6 (a) we have represented on one hand the daily event related transaction share (black line) and the daily standard deviation of the location baseline-event footprint difference distribution $\sigma(|EB(L, \triangle t) - EF(L, E, \triangle t)|)$. Fig. 6 (a) shows the resulting emotional impact metric taken day by day for both events in all three locations. It's remarkable how the emotional impact fades progressively out during the week in which the incident happened in both cases, but also how both events affect the chosen locations in different ways: we could say that both deaths have similar impact in Chelsea, but while Edinburgh has been definitively more impacted by Mandela's, Walker's death left a deeper mark in Manchester. The results obtained for Edinburgh shows us the role of the event's share in the impact metric: while Walker's death presents a more diverging emotional footprint from the emotional baseline of the Scottish capital, the share is much lower than Mandela's and therefore the overall impact.

5.3. The named mood emotional impact

As we explained in the Subsection 3.3, the emotional impact can be expressed by how particular emotional states gain or lose share. In Fig. 7 we wanted to first show a direct comparison of both events over all gathered transactions in all three locations; we see for example that while Mandelas death *impressed* more people, Walkers death left more people *discontented* and *expectant*.

In Fig. 8 we have plotted the change in the top 15 named moods that have been impacted the most by both events in the three cities. In general, we observe that typically strong emotional baseline named moods, such as *longing*, *attentive* or *helpful* are highly impacted in terms of share loss by both events, while named-moods on the other side of the y-axis (negative valences) profit from this loss.

Mandelas death massively *impressed* people in all three locations. *Expectancy* was also observable in Manchester and Edinburgh, while Chelsea reacted more *contemplatively*. Remarkable uplift of *apathetic* feelings and people *taken aback* in all locations.

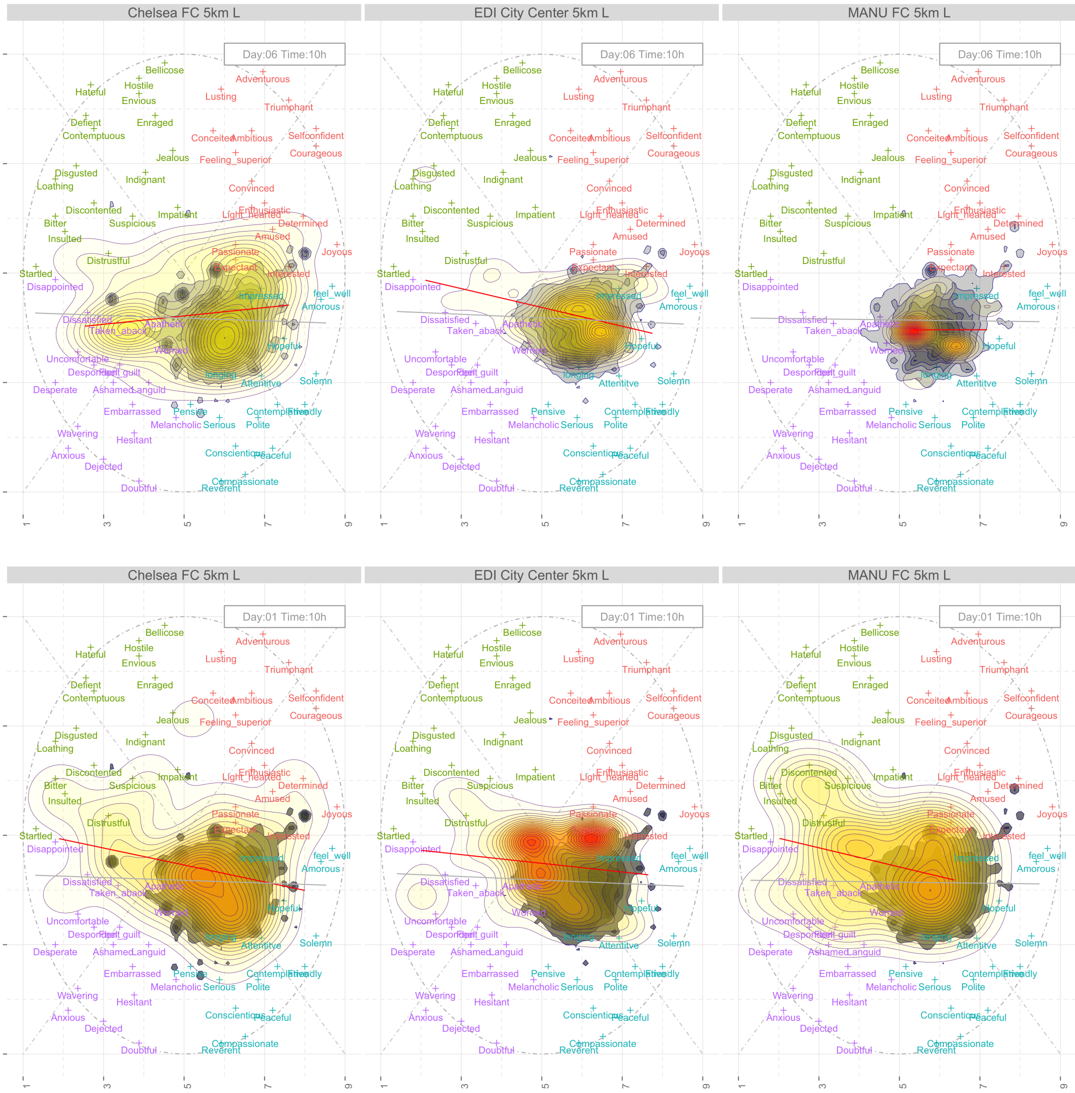

Figure 4: Emotional Impact for Mandela's death and Walker's Death in 3 Locations at Named Mood Level

Walkers decease released a generalized *discontent* in the English cities. *Apathy* is also noticeable in a general note as well as *expectancy*. Manchester and Edinburgh show an increase of *distrust* and *discomfort*, while in Chelsea people also feel *suspicious* and *insulted*.

As we have seen, with our method we can precisely say how much the three different locations have been impacted by both events, but also we can qualify this impact in terms of particular moods.

6. Conclusions

In this paper we present a new approach to quantify the emotional impact of an event on a physical location based on the analysis of Social Media interactions that have been geo-located in this location.

To achieved that, we first introduced the concepts of emotional baseline for a location and emotional footprint for an event based on the analysis of user generated content posted over SM in the place under analysis. After that,

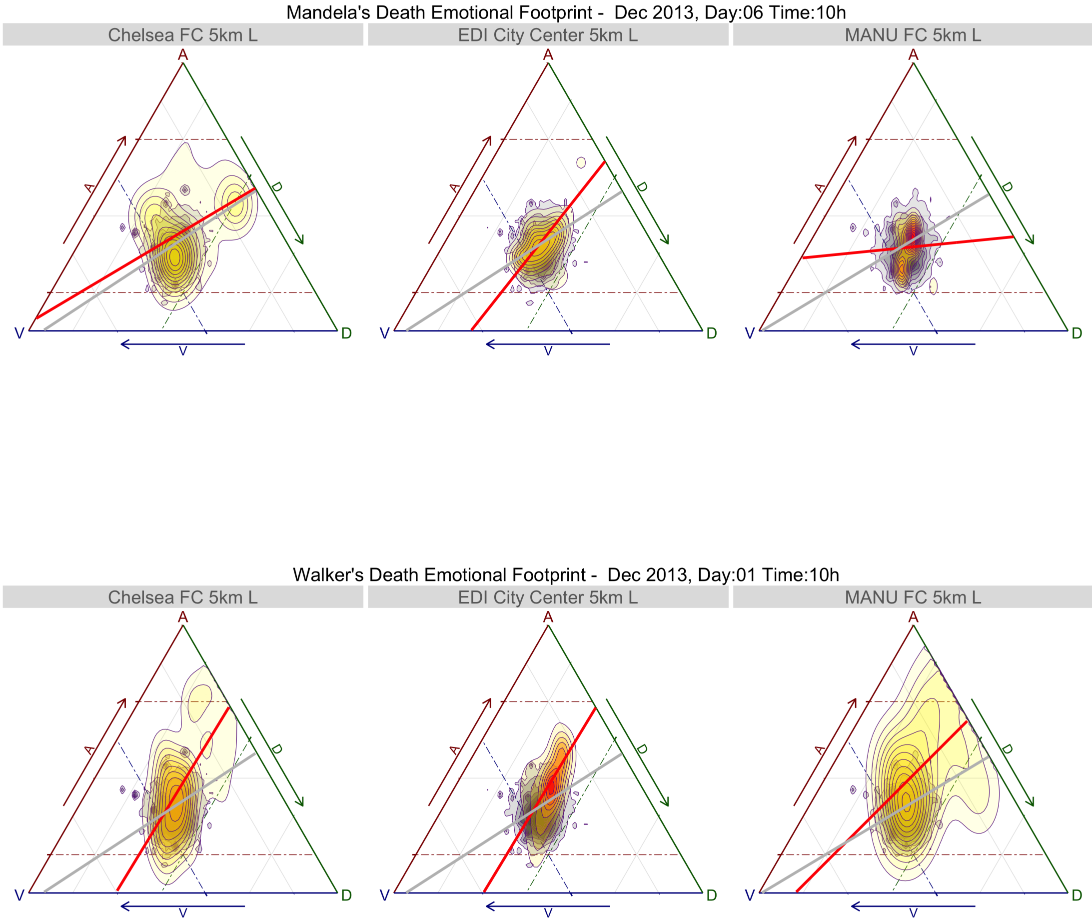

Figure 5: Emotional Impact for Mandela's death and Walker's Death in 3 Locations at Named Mood Level

we defined the emotional impact as the difference between both concepts and provided a mechanism to measure this impact at a much finer granular level, namely for each particular existing mood.

Our method builds upon following components: a) the well-established (P)leasantness or (V)alence-(A)rousal-(D)ominance emotional state model introduced by Russell, to model emotions, b) an extended version of the Affective Norms for English Words, to extract emotions from the Social Media user generated content and c) and an evolution of the Russell's circumplex model to map the $[v, a]$ scores to one of the set of named emotional states derived from, such as *Impatient, Hopeful, Amorous, etc*.

Both the emotional baseline of a location and the emotional footprint of an event are defined by the multivariate kernel density function applied to the whole set of $[v, a, d]$ scores gathered over the define time period. Our method works in the bi-dimensional $[v, a]$ space to enable the mapping to named moods on one hand, and in the $[v, a, d]$ three-dimensional space to consider the effect of the dominance component on the other hand.

To evaluate our approach, we implement a system based on Twitter and discuss the results in different scenarios for three known locations in Great Britain: Edinburgh city center, Chelsea in London and the surroundings of Old

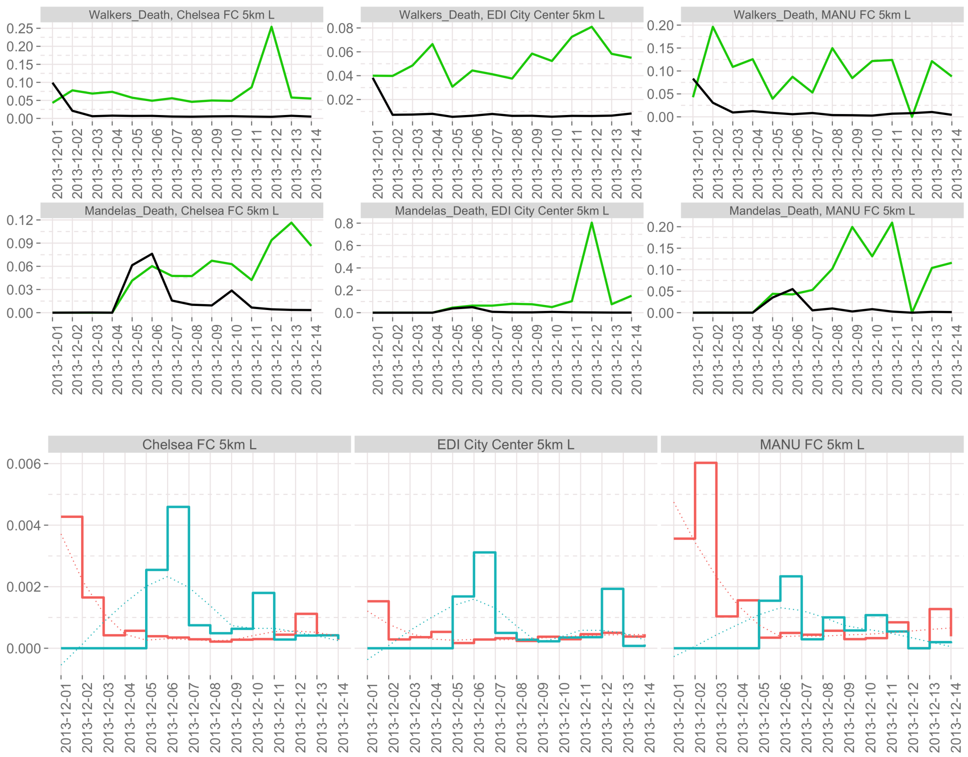

Figure 6: (a) Share -black- and difference σ -green- for the locations and the events (b) Impact metrics for Mandela's death -cyan- and Walker's death -red- for the locations over time

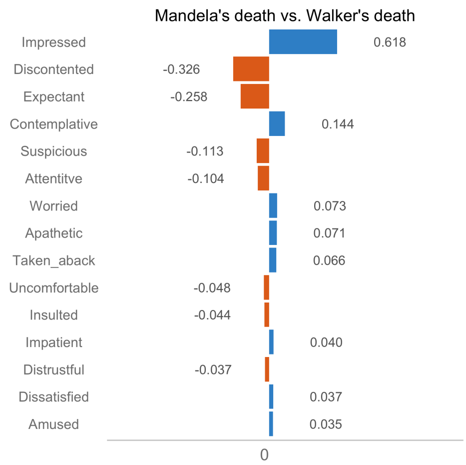

Figure 7: Mandela's death vs. Walker's death named moods differences

Trafford in Manchester. Our analysis focused on quantifying the emotional impact of Nelson Mandela's death and Paul Walker's decease at the beginning of Dec. 2013, which we have carried out with different granularity levels –hourly, daily and bi-weekly– showing in a very thorough manner the performance of our method and uncovering the potential to apply it in real-world applications.

The applications of emotional profiling of locations in general and emotional impact measuring in particular are

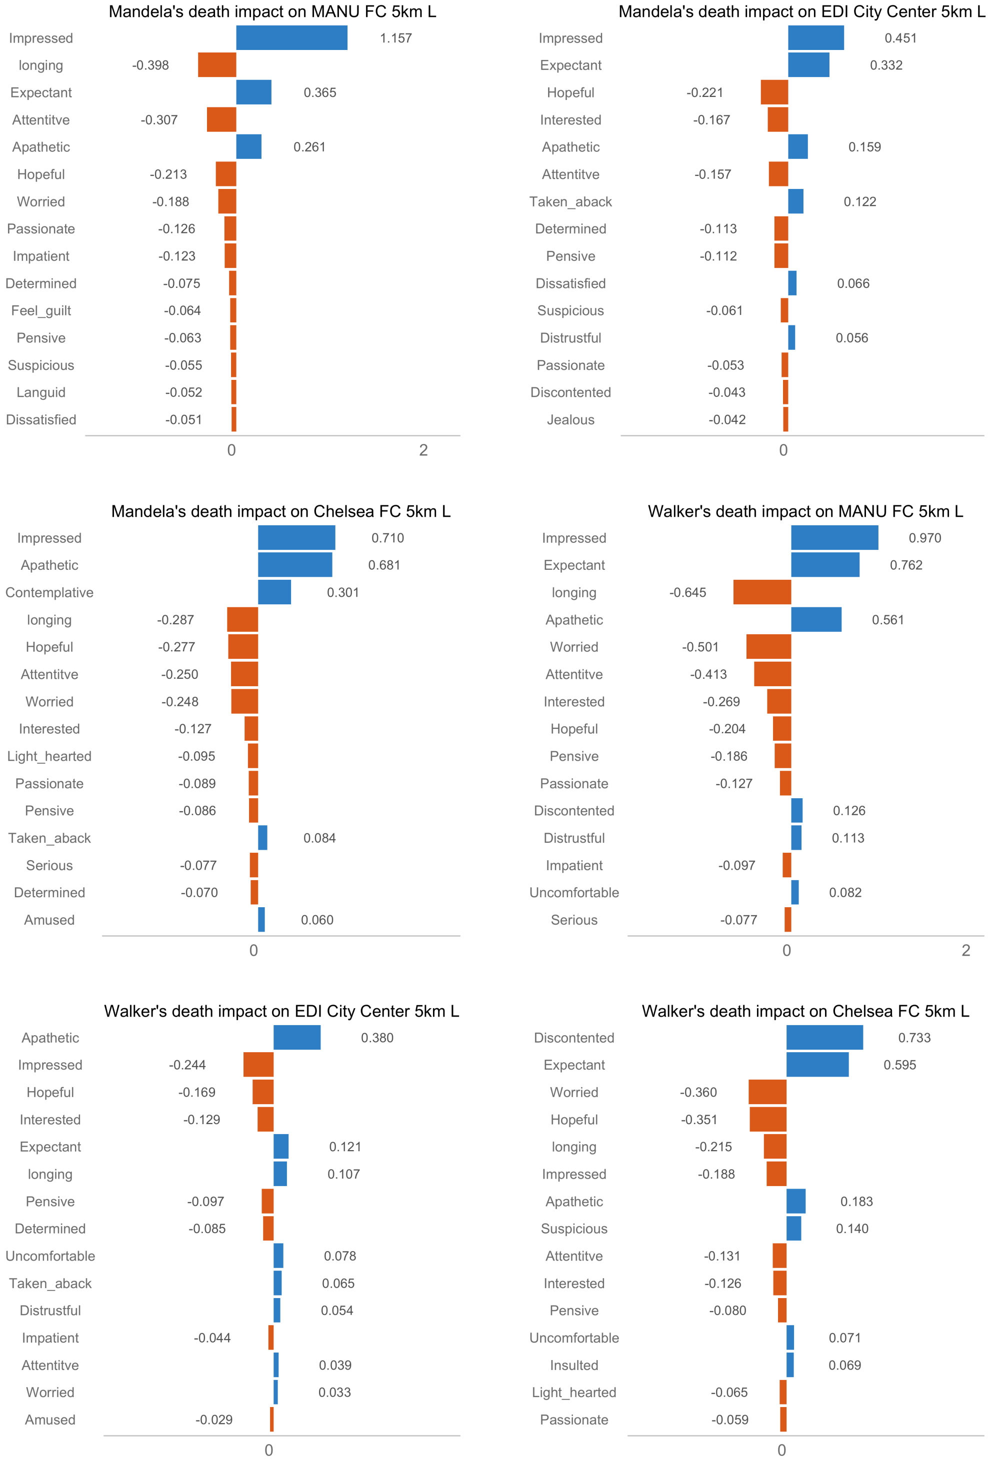

Figure 8: Emotional Impact for Mandela's death and Walker's Death in 3 Locations at Named Mood Level

countless. This kind of insights open a new door to marketing activities (e.g.: choosing the right marketing message that fits best the emotional baseline of a location, identifying the best set of promotional activities based on emotional impact, etc.), tailoring of political campaigns (e.g.: selecting the right wording in the massages and measure the outcome) or at a particular level, even finding the right place to live based on the emotional profile of the potential neighbours and their emotional reaction to events. These are just a few examples of the countless applications of the output of this piece of work.

To continue the research initiated in this paper, we suggest exploring the adoption of a user centric approach – for example, creating emotional profiles of users over a longer period of time, that then are mapped to locations for better consistency or considering the segmentation by gender and educational class already present in the extended ANEW –. Another interesting area would be developing approaches for removing the digital bias to make the insights representative for the entire population of a location, not just the geo-located SM users. With enough SM history, understanding emotional profile changes in locations over time or even clustering events depending on their emotional profile would massively enrich this research line as well.

Acknowledgements

This paper has been developed with the financing of projects TIC-5991 and TIN2013-40658-P.

References

[1] Z. Cheng, J. Caverlee, K. Lee, You are where you tweet: a content-based approach to geo-locating twitter users, in: Proceedings of the 19th ACM international conference on Information and knowledge management, ACM, 2010, pp. 759–768.

[2] E. Cho, S. A. Myers, J. Leskovec, Friendship and mobility: user movement in location-based social networks, in: Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM, 2011, pp. 1082–1090.

[3] C. Nuernbergk, Follow-up communication in the blogosphere: A comparative study of bloggers linking to professional and participatory media, Digital Journalism 2 (3) (2014) 434–445.

[4] J. Macnamara, A. Zerfass, Social media communication in organizations: The challenges of balancing openness, strategy, and management, International Journal of Strategic Communication 6 (4) (2012) 287–308.

[5] C. Porcel, A. Tejeda-Lorente, M. A. Martínez, E. Herrera-Viedma, A hybrid recommender system for the selective dissemination of research resources in a technology transfer office, Inf. Sci. 184 (1) (2012) 1–19.

[6] Á. Tejeda-Lorente, C. Porcel, E. Peis, R. Sanz, E. Herrera-Viedma, A quality based recommender system to disseminate information in a university digital library, Information Sciences 261 (2014) 52–69.

[7] J. Serrano-Guerrero, J. A. Olivas, F. P. Romero, E. Herrera-Viedma, Sentiment analysis: A review and comparative analysis of web service, Information Science 311 (2015) 18–38.

[8] B. Pang, L. Lee, S. Vaithyanathan, Thumbs up?: sentiment classification using machine learning techniques, in: Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10, Association for Computational Linguistics, 2002, pp. 79–86.

[9] M. Wollmer, B. Schuller, F. Eyben, G. Rigoll, Combining long short-term memory and dynamic bayesian networks for incremental emotion-sensitive artificial listening, Selected Topics in Signal Processing, IEEE Journal of 4 (5) (2010) 867–881.

[10] R. Wicentowski, M. R. Sydes, emotion detection in suicide notes using maximum entropy classification, Biomedical informatics insights 5 (Suppl 1) (2012) 51.

[11] C. Havasi, E. Cambria, B. Schuller, B. Liu, H. Wang, Knowledge-based approaches to concept-level sentiment analysis, IEEE Intelligent Systems 28 (2) (2013) 0012–14.

[12] S. Poria, A. Gelbukh, E. Cambria, A. Hussain, G.-B. Huang, Emosenticspace: A novel framework for affective common-sense reasoning, Knowledge-Based Systems 69 (2014) 108–123.

[13] J.-M. López-Gil, R. Gil, R. García, C. A. Collazos, Emotionsonto: an ontology for developing affective applications, Journal of Universal Computer Science 20 (13) (2014) 1813–1828.

[14] M. Bradley, P. Lang, Affective norms for english words (anew): Technical manual and affective ratings, Gainsville, FL: Center for Research in Psychophysiology, University of Florida.

[15] M. M. Bradley, P. J. Lang, Measuring emotion: the self-assessment manikin and the semantic differential, Journal of behavior therapy and experimental psychiatry 25 (1) (1994) 49–59.

[16] A. Mehrabian, Basic dimensions for a general psychological theory: Implications for personality, social, environmental, and developmental studies, Oelgeschlager, Gunn & Hain, Incorporated, 1980.

[17] J. A. Russell, A circumplex model of affect., Journal of personality and social psychology 39 (6) (1980) 1161.

[18] W. Revelle, K. R. Scherer, Personality and emotion, Oxford companion to emotion and the affective sciences (2009) 304–306.

[19] J. Bernabé-Moreno, A. Tejeda-Lorente, C. Porcel, H. Fujita, E. Herrera-Viedma, Caresome: A system to enrich marketing customers acquisition and retention campaigns using social media information, Knowledge-Based Systems.

[20] E. Hudlicka, What are we modeling when we model emotion?, in: AAAI spring symposium: emotion, personality, and social behavior, 2008, pp. 52–59.

[21] A. Mehrabian, J. A. Russell, An approach to environmental psychology., the MIT Press, 1974.

[22] R. E. Thayer, The biopsychology of mood and arousal, Oxford University Press, 1989.

[23] D. Watson, A. Tellegen, Toward a consensual structure of mood., Psychological bulletin 98 (2) (1985) 219.

[24] P. J. Lang, M. M. Bradley, B. N. Cuthbert, Emotion, motivation, and anxiety: brain mechanisms and psychophysiology, Biological psychiatry 44 (12) (1998) 1248–1263.

[25] A. B. Warriner, V. Kuperman, M. Brysbaert, Norms of valence, arousal, and dominance for 13,915 english lemmas, Behavior research methods 45 (4) (2013) 1191–1207.

[26] M. L. Võ, M. Conrad, L. Kuchinke, K. Urton, M. J. Hofmann, A. M. Jacobs, The berlin affective word list reloaded (bawl-r), Behavior research methods 41 (2) (2009) 534–538.

[27] J. Redondo, I. Fraga, I. Padrón, M. Comesaña, The spanish adaptation of anew (affective norms for english words), Behavior research methods 39 (3) (2007) 600–605.

[28] A. P. Soares, M. Comesaña, A. P. Pinheiro, A. Simões, C. S. Frade, The adaptation of the affective norms for english words (anew) for european portuguese, Behavior research methods 44 (1) (2012) 256–269.

[29] W.-L. Wei, C.-H. Wu, J.-C. Lin, A regression approach to affective rating of chinese words from anew, in: Affective Computing and Intelligent Interaction, Springer, 2011, pp. 121–131.

[30] C. Strapparava, A. Valitutti, et al., Wordnet affect: an affective extension of wordnet., in: LREC, Vol. 4, 2004, pp. 1083–1086.

[31] S. Ramaswamy, C. G. Healey, Visualization of the sentiment of tweets, in: SAS Analytics 2011 Conference, SAS, 2011.

[32] S. Gobron, J. Ahn, G. Paltoglou, M. Thelwall, D. Thalmann, From sentence to emotion: a real-time three-dimensional graphics metaphor of emotions extracted from text, The Visual Computer 26 (6-8) (2010) 505–519.

[33] S. Stieglitz, L. Dang-Xuan, Emotions and information diffusion in social mediasentiment of microblogs and sharing behavior, Journal of Management Information Systems 29 (4) (2013) 217–248.

[34] A. Chmiel, P. Sobkowicz, J. Sienkiewicz, G. Paltoglou, K. Buckley, M. Thelwall, J. A. Hołyst, Negative emotions boost user activity at bbc forum, Physica A: statistical mechanics and its applications 390 (16) (2011) 2936–2944.

[35] G. Paltoglou, M. Thelwall, Seeing stars of valence and arousal in blog posts, Affective Computing, IEEE Transactions on 4 (1) (2013) 116–123.

[36] P. S. Dodds, K. D. Harris, I. M. Kloumann, C. A. Bliss, C. M. Danforth, Temporal patterns of happiness and information in a global social network: Hedonometrics and twitter, PloS one 6 (12) (2011) e26752.

[37] J. Bernabé-Moreno, A. Tejeda-Lorente, C. Porcel, E. Herrera-Viedma, A new model to quantify the impact of a topic in a location over time with social media, Expert Systems with Applications 42 (7) (2015) 3381–3395.

[38] J. Bernabé-Moreno, A. Tejeda-Lorente, C. Porcel, E. Herrera-Viedma, Leveraging localized social media insights for early warning systems, Procedia Computer Science 31 (2014) 1051–1060.

[39] T. Sakaki, M. Okazaki, Y. Matsuo, Earthquake shakes twitter users: real-time event detection by social sensors, in: Proceedings of the 19th international conference on World wide web, ACM, 2010, pp. 851–860.

[40] A. Chatfield, U. Brajawidagda, Twitter tsunami early warning network: a social network analysis of twitter information flows.

[41] D. Yates, S. Paquette, Emergency knowledge management and social media technologies: A case study of the 2010 haitian earthquake, International Journal of Information Management 31 (1) (2011) 6–13.

[42] K. R. Scherer, What are emotions? and how can they be measured?, Social science information 44 (4) (2005) 695–729.

[43] J. Broekens, In defense of dominance: Pad usage in computational representations of affect, International Journal of Synthetic Emotions (IJSE) 3 (1) (2012) 33–42.

[44] E. Parzen, On estimation of a probability density function and mode, The annals of mathematical statistics (1962) 1065–1076.

[45] W. N. Venables, B. D. Ripley, Modern applied statistics with S, Springer Science & Business Media, 2002.

[46] M. Oussalah, F. Bhat, K. Challis, T. Schnier, A software architecture for twitter collection, search and geolocation services, Knowledge-Based Systems 37 (2013) 105–120.

[47] B. O'Connor, M. Krieger, D. Ahn, Tweetmotif: Exploratory search and topic summarization for twitter., in: ICWSM, 2010.

[48] T. Kiss, J. Strunk, Unsupervised multilingual sentence boundary detection, Comput. Linguist. 32 (4) (2006) 485–525.

[49] Q. X. Yang, S. S. Yuan, L. Zhao, L. Chun, S. Peng, Faster algorithm of string comparison, Pattern Analysis & Applications 6 (2) (2003) 122–133.

[50] E. Clark, T. Roberts, K. Araki, Towards a pre-processing system for casual english annotated with linguistic and cultural information, in: Proceedings of the Fifth IASTED International Conference, Vol. 711, 2010, pp. 044–84.

2. Geo-localized Campaigning and Quality of Service Monitoring

The journal papers associated to this part are:

2.1 CARESOME: A System to Enrich Marketing Customers Acquisition and Retention Campaigns Using Social Media Information

- Bernabé-Moreno, J., A. Tejeda-Lorente, C. Porcel, H. Fujita, and E. Herrera-Viedma. "CARESOME: A system to enrich marketing customers acquisition and retention campaigns using social media information." Knowledge-Based Systems 80 (2015): 163-179.
  - Status: **Published**.
  - Impact Factor (JCR 2013): 1.977
  - Subject Category: Computer Science, Artificial Intelligence. Ranking 17 / 688 (**Q1**).
  - Subject Category: Computer Science, Information Systems and Management. Ranking 10 / 122 (**Q1**).
  - Subject Category: Computer Science, Management Information Systems. Ranking 6 / 76 (**Q1**).
  - Subject Category: Computer Science, Software. Ranking 32 / 1193 (**Q1**).

CARESOME: A System to Enrich Marketing Customers Acquisition and Retention Campaigns Using Social Media Information

J. Bernabé-Moreno[a,*], A. Tejeda-Lorente[a], C. Porcel[b,*], H. Fujita[c], E. Herrera-Viedma[a,*]

[a]*Department of Computer Science and A.I., University of Granada, Granada E-18071, Spain*
[b]*Department of Computer Science, University of Jaén, Jaén E-23071, Spain*
[c]*Iwate Prefectural University (IPU), Iwate, Japan*

Abstract

The enabling of geo-localization for Social Media content opens the door to a new set of applications based on the voice of the customer. For any company it is critical to understand both their own and their competitors' strengths and weaknesses in all locations where they offer a service. With this motivation we created a Customers Acquisition and REtention system based on SOcial MEdia (CARESOME). Our system extracts and separates all social media interactions in a given location by market player and communication purpose and quantifies the impact of each single interaction over a given time period. To model the impact of the social media interactions, CARESOME relies on a set of metrics based on both intrinsic and extrinsic components —including Entity Engagement Index, Differential Perception Factor, Tie-Strength and Number of Exposed users—. In addition to the definition of our impact quantification metrics, we provide a thorough discussion about the design decisions taken to build our system. To illustrate the behaviour of our system, we show-case a real world scenario from the airline industry based on two major airports in Great Britain.

Keywords:
Intrinsic Impact; Extrinsic Impact; Social CRM; Customers Retention;Customers Acquisition; Localized Social Media; Ubiquitous Insights

1. Introduction

Social Media (SM) started as a space where anybody with an account could interact with any other user, share content, express their own personal views, etc. without being subjected to any kind of censorship. As a side effect of this democratization of the Web, the relationship between a company and its customers and stakeholders went through an unprecedented transformation [1]. For the first time, customers could engage in a near real time manner with companies and brands [2]. The advent of SM radically changed the way customers engage with service providers or product vendors. Any customer could express in an unfiltered way his/her opinion about a brand, a service, a price increase, etc. and the result of it was publicly available in a near real time manner to other customers or customers-to-be. The *killer application* of SM in the consumer market has been the customer empowerment. The customer feedback, that used to be trapped in the traditional offline *word-of-mouth* modus operandi, is now available to each and every user willing to know more about the quality of service of any company in the world. SM made these communication barriers fall and changed the customer-company engagement rules for ever, as different types of business are using customer data for better comprehension on customers data [3].

Companies had been left with no choice but embracing the new customers' engagement channel and developing customers' acquisition and customers' retention strategies on top: leaning back and not doing anything was no longer an option [4]. Early adopters managed to build up a new form of competitive advantage relying on both own

*Corresponding author. Tel.: +34-958-24-4258 ; fax: +34-958-24-3317 .
E-mail address: juan.bernabe-moreno@webentity.de,cporcel@ujaen.es,viedma@decsai.ugr.es.

customers' binding and competitors' customers caption strategies. Spotting signs of own customers' satisfaction decay in a given location made companies trigger local customers' retention campaign, for example in [5], stating that purchasing behaviors can be significance altered by supplying consumers with seller-generated information. Likewise, localized satisfaction decay in competitors' customers base pushed companies to implement highly aggressive regional campaigns to take advantage of the weak spot by mitigating diffusion [5].

The need for geo-localized systems to monitor the customer satisfaction at a local scale and to assess the impact of customers' interactions with the brand over SM, emerged [6]. Early warning systems —the equivalent in other domains— have been increasingly adopted in the field of disaster prevention as the sensorial technique allowed for semi-automatic monitoring. There are countless applications for early detection of earthquakes [7, 8, 9], pandemics [10, 11], flood and other natural hazards [12]. In the financial domain fast alerting system have been employed for a wide range of purposes: for example, all variety of economic indicators have been used at a macro level to assess the vulnerability of emerging and existing markets [13, 14, 15] and to detect financial crisis in their early stages [16, 17], but also at a much more micro level to detect for example critical transactions [18], etc.

When the access to the world wide web (WWW) escaped the desktop boundaries and became mobile and pervasive —mainly because of WiFi and the third generation mobile cellular system for networks based on the GSM standard, Universal Mobile Telecommunications System (UMTS)—, the SM platform providers leveraged the geo-location of the interactions as a *highly enriching* additional information source. There was such a demand for location enriched user interactions, that new platforms less focused on content but more focused on location, like Foursquare emerged and quickly started conquering the market. Well established content and community based platforms traditionally positioned as the medium where customers engaged with brands (such as Facebook or Twitter) immediately reacted enabling the localization of the user interactions to improve the user experience. The geo-localization of customers' interactions opened a new door for companies to better understand their own customers' base and develop strategies to take customers away from competitors in their own favor. Understanding the impact of each and every interaction over SM on one hand and triggering on time the appropriate reaction proved to be two essential factors for succeeding in a customers retention or customers acquisition strategy [19]. Even if companies are heavily investing in standing up SM care teams, the SM adoption makes the handling of each and every SM interaction far from scalable, which introduces the need for a way of quantifying their impact.

In this paper we present our Customer Acquisition and REtention system over SOcial MEdia, which leveraging Big Data technologies [20], implements a framework based on geo-localized Tweets, to measure the impact of the interactions created in a location on a brand or institution, or any kind of entity. Our system takes a holistic view over all factors that play a role in the impact perception within a SM context, such as authors engagement with the entity, followers exposure to the interactions, Tie-Strength between authors and followers, etc. As outcome, CARESOME produces a set of metrics for a location over a given time frame aggregated by communication purpose category to enable the response by different company departments (e.g.: customer care is likely to focus on the categories criticism and complaints, while marketing would rather be interested in measuring the impact of a new campaign based on positive feedback , etc). Additionally, the result of these metrics is packed into impact categories to enable faster decision making, as time to reaction is proven to be a critical success factor in every early-warning like system. In other words, CARESOME turns the information extracted from the different SM channels into actionable insights for companies to steer their customer acquisition and loyalization campaigns based on opinions of customers.

We started our work presenting all the background our research is built upon (Section 2). In Section 3, we introduced the impact quantifying framework and define all relevant metrics. Section 4 explains the CARESOME system architecture and Section 5 presents a real-world scenario and provides a discussion about the system performance and our design decisions. Section 6 closures this work pointing out future research lines and summarizing our conclusions.

2. Background

Nowadays, almost every company relies on SM as a communication channel to push company messages and offers, but also increasingly to obtain unfiltered feedback from both existing and prospective customers. Many studies have focused on different aspects of the SM adoption: Kaplan et al. [21] highlighted the need for the integration of SM with traditional media to reach customers more efficiently, while defending the advantages of SM to engage with customers in a time-close and high-efficient manner. Mangold et al. [22] built upon the idea of considering SM as integral part of the promotion mix, emphasizing the benefit of a less controlled environment to better understand customers.

Several papers focused on researching the role of SM in business and corporations. Jansen et al. in [23] analyzed the corporate image impact of all interactions related to a brand created over the Twitter channel. In [24], Li et al. explained the positive impact of the user engagement over the Twitter company channels on the corporate reputation. In [25] Java et al. demonstrated how similar intentions foster connectivity between users and community building around brands and institutions. Plenty of studies shed light on how companies shall deal with SM related issues like trust and distrust within online communities [26] and protection of user's information[27].

SM rapidly moved from being *yet another channel* in the communication strategy of a company to be labeled as a *game changer* to engaging with customers: Hennig et al. [1] explained how microblogging was shaking traditional business models by increasing the role of product quality, especially reducing the time window where product new adopters didn't have any feedback on the product. Culnan et al. [28] pointed out the need for brands to create communities to exploit the full potential of the virtual customer environments. In [29] the link between SM engagement and profitability of online companies was analysed by Chan and his co-authors. In [30], Rapp and his co-authors analysed the role of SM from the seller, retailer and consumer perspective, demonstrating the value of the SM interactions for better conversion rate.

The effect of the Worth-of-Mouth (WoM) marketing has been extensively researched together in the SM context. Chevalier and his co-authors analysed in [31] the effect of book reviews. Villanueva et al. [32] researched the differences in terms of loyalty and equity of customers being acquired through marketing-induced activities vs. WoM gained customers, pointing out performance differences. Bolton established back in 1998 [33] a modeled based on the link between customers retention and customers satisfaction and Rishika et al. [34] empirically proved the effect of increased SM engagement on the customer visit frequency and customer value.

As proved in [35] and [36], the spreading of bad news takes place really fast over the SM channel, which corroborates their value for the promptly detection of customers' complaints, service outages, etc. Countless papers built upon the fast news spreading aspect of SM: in [37], 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 in [38] a method to detect epidemic expansion on early stages. In [39], Middleton and his co-authors present a near real time system to map crisis based on several geo-localization techniques of SM information. In the same research line, Yin et al. in [40] present a system that implements text mining and natural language processing (NLP) techniques to extract situation awareness information from Twitter to support crisis coordination and emergency response. In [41] Colbaugh and Glass employed a stochastic model for dynamic of the interactions based on the underlying network structure to generate useful predictions about the spread of information. The US Homeland department pioneered the usage of SM to collect real time information about incidents, quantify their extent, monitor their evolution and channel the proper response —programme SMART-C (SM Alerts and Response to Threats to Citizens)—[42].

Predicting (i.e., customers) behaviours in SM for management decision making is still challenging tasks [43, 44, 45]. The analysis of SM content and engagement to predict upcoming events has been also intensively researched. In [46] the Bothos, Apostolou and Mentzas explain how agents constantly analysing social media content according to the Belief-Desire-Intentions paradigm can extract enough sentiments and assessments to enable informed decision making in the markets they operate.

Our Impact metrics, as we are going to explain later in this paper, relies on how influential a particular SM's user is. Modeling influence in SM channels has been subject of intense research over the last few years. Kwak [47] defined 3 metrics aimed at quantifying the *social influence*: the so called *propagation influence*, based on the Google Search PageRank algorithm [48], *followers influence* —more followers implies more influence—, and *re-tweet influence* —more re-tweets means more influence—. Ye and Wu [49] 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—. Cha [50] 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. [51] develop a mechanism to quantify how the exposure to other users is making them adopt a new behavior. Yang et al. in [52] add a new dimension to the influence computing, namely the response immediacy in their influence modeling for an online health care community.

There have been several studies showing how the Tie-Strength between two SM users plays an important role in the perception of SM interactions. Marsden and his co-authors in [53] back in 1984 laid the foundations for measuring the Tie-Strength after Mark Granovetter introduced the concept in 1973 in his paper "The Strength of Weak Ties" [54]. In [55], a model to predict tie strength by exploiting social media interaction parameters is discussed. The work done by Haythornthwaite [56] confirmed that more strongly tied pairs communicate more frequently, maintain more and different kinds of relations and use more media to communicate. Grabowicz et al. in [57] analyzed the relationship between SM links and real-world tie-strength and Pan et al. in [58] attempted to quantify the role of tie-strength plays in scientific collaboration networks. Shin et al. presented a method to quantify the degree of user sociability in SM relying on the tie-strength[59].

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 for decision making. The Impact of a SM interaction with a brand can be modeled from two perspectives: intrinsically —reflecting the impact perception on the SM interaction author— and extrinsically —which captures how the SM interaction impacted its author's SM network—.

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. CARESOME provides also the flexibility of defining own categories in the event that the standard ones are not suitable.

3.1. Preliminary definitions

Before starting with the definition of our framework, a set of concepts to support our metrics needs to be established:

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

$$U \equiv \{u_i\}\ (i = 1, ..., n),\ InLocation(u_i, L, \triangle t) \tag{1}$$

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

$$SN\ (u_i) \equiv \{u_j\}(j = 1, ..., n),\ \forall\ u_j \in SN(u_i),\ Follows(u_i, u_j) \tag{2}$$

$Follows(u_i, u_j)$ is a relation 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 $Follows(u_i, u_j) \not\Leftrightarrow Follows(u_j, u_i)$

The fact that a user u_j is part of the SN of another user u_i does not necessarily mean that u_j has to be located in the same location L of user u_i: $u_j \in SN(u_i), u_i \in U \not\Rightarrow u_j \in U$, as $u_j \in SN(u_i), u_i \in U \not\Rightarrow InLocation(u_j, L, \triangle t)$

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

$$SN(U) \equiv \cup_{i=1}^{|U|} SN(u_i) \tag{3}$$

Definition 4. *A Social Media Interaction it represents the atomic piece of content authored by the user u_i during the time $\triangle t$ in a Social Media Platform (e.g.: a tweet, a re-tweet).*

The function $Author(u_i, it_i, \triangle t)$ returns $True$ if u_i created the interaction it_i in the time period $\triangle t$, and $False$ otherwise.

The time interval $\triangle t$ 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. *We define all user interactions* (*Interactions*) *for a given user u_i over a time interval $\triangle t$, as:*

$$Interactions\ (u_i, \triangle t) \equiv \{it_i\}(i = 1, ..., n),\ \forall it_i \in Interactions\ (u_i, \triangle t),\ Author(u_i, it_i, \triangle t) \tag{4}$$

Definition 6. *The set of User Foreign Interactions ($ForeignInteractions(u_i, \triangle t)$) represents all Interactions with a direct mention to the user u_i but not authored by him/her:*

$$ForeignInteractions(u_i, \triangle t) \equiv \{it_j\}\forall it_j \in ForeignInteractions(u_i, \triangle t), \neg Author(u_i, it_j, \triangle t), DirectMentioned(it_j, u_i) \tag{5}$$

$DirectMentioned(it_i, u_i)$ is a function retrieving $True$ if the user u_i is explicitly mentioned in it_i. In Twitter, the User Foreign Interactions include re-tweets, mentioned and replies.

$ForeignInteractions(u_i, \triangle t) \cap Interactions(u_i, \triangle t) = \oslash$

Let's illustrate it with one example; for a user with the user name *@user1*, a tweet created by *@user2* saying "*Happy birthday @user1"* represents a foreign interaction for *@user1*. If @user3 retweets it, it counts as well as a foreign transaction for *@user1*, who is mentioned in the tweet text, but also as a foreign interaction for *@user2*, as his/her tweet has been re-tweeted.

Definition 7. *The set of Direct Mention Interactions is as a subset of Interactions $(u_i, \triangle t)$ defined as follows:*

$$DirectMentionInteractions\ (u_i, \triangle t) \equiv \{it\},\ \forall it_i \in Interactions\ (u_i, \triangle t), \exists u_j \mid it_i \in ForeignInteractions\ (u_j, \triangle t) \tag{6}$$

Intuitively, DirectMentionInteractions represents all the interactions created by the user u_i where any other user is explicitly mentioned. Obviously, $DirectMentionInteractions\ (u_i, \triangle t) \subseteq Interactions\ (u_i, \triangle t)$

Definition 8. *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 9. *We define the set of Interactions for a given user u_i with the entity E over a time interval $\triangle t$ as:*

$$Interactions\ (u_i, E, \triangle t) \equiv \{it\},\ \forall it_i \in Interactions\ (u_i, E, \triangle t),\ Author(u_i, it_i, \triangle t)\ \wedge related(it_i, E) \tag{7}$$

Where $related(it_i, E)$ is a NLP membership function retrieving $True$ if the interation 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:

$$Engaged\ (u_i, E, \triangle t) \equiv True, \exists it_i,\ it_i \in Interactions\ (u_i, E, \triangle t), u_i \in U \cup SN(U) \tag{8}$$

Where u_i is the user, E is the representation of the Entity, $\triangle t$ is the time span specified consisting of two components ($t\_startdate$ and $t\_enddate$), it_i represents a social media interaction and $Interactions\ (u_i, E, \triangle t)$ represents the interactions of the user u_i related to the Entity E in the time interval $\triangle t$, 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, the so called *Entity Engagement Index (EEI)*:

$$EEI(u_i, E, \triangle t) = \frac{|Interactions\ (u_i, E, \triangle t)|}{|\cup_{k=1}^{|E|} Interactions\ (u_i, E_k, \triangle t)|} \tag{9}$$

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

The Entity Engagement Index can also be expressed as a share of the interactions related to one entity over all interactions:

$$EEI(u_i, E, \triangle t) = \frac{|Interactions\ (u_i, E, \triangle t)|}{|Interactions(u_i, \triangle t)|} \tag{10}$$

3.3. Social Media Communication Intent

Behind each and every posts or tweet or, in general, piece of content authored by a user in a Social Media platform there is an underlying communicative 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 10. *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 a particular communicative intention (such as praise, criticism, information inquiry, complaints, etc).*

Even if the boundaries might not be crisp, we can assign each interaction to a *leading Purpose Category* within the set of purpose categories considered PC:

$$\forall it_i \in Interactions\ (u_i, E, \triangle t), \exists p_k, Purpose(it_i) = p_k, p_k \in PC \tag{11}$$

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

$Interactions(u_i, P, \triangle t)$ represents the set of all interactions authored by a user u_i over the period of time $\triangle t$ whose leading Purpose Category is P.

3.4. Differential Perception Factor, Exposure and Tie Strength

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

We introduce the so called *Differential Perception Factor* modelled as *Purpose Share* (see the previous Def. 10), which allows for latterly defining a correction factor to remove the SM behavioural bias:

$$DPF(u_i, E, P, \triangle t) = \frac{|Interactions\ (u_i, E, \triangle t) \cap Interactions\ (u_i, P, \triangle t)|}{|Interactions\ (u_i, P, \triangle t)|} \tag{12}$$

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.

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

$$Exposed\ (u_i, u_j, E, \triangle t) = \begin{cases} True, & u_j \in SN(u_i), \exists it_k, it_k \in Interactions(u_i, E, \triangle t), P(read(u_j, it_k, \triangle t)) \geq Threshold \\ False, & otherwise \end{cases} \tag{13}$$

where $P(read(u_j, it_k, \triangle t))$ is the probability that the user u_j reads the content posted in the interaction it_j in the designated time $\triangle t$. 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:

$$ExposedUsers(u_i, E, \triangle t)\ \equiv \{u\},\ \forall u_j,\ Exposed(u_i, u_j, E, \triangle t)\ =\ True,\ u_i \in U \tag{14}$$

An additional yet quite relevant aspect we incorporate to the Impact definition is the relationship between the author of the social media interaction and the user in his/her SM network. Depending on this relationship, the level of perceived relevance might vary. For example, if a given user u_i is a good friend of $u_j, u_j \in SN(u_i)$, the relevance, the u_j perceives u_j's posts to have is higher than it would be if there was practically no link between these users apart from the fact that u_j is part of the SM network of u_i. Thus, we define Tie-Strength between two social media users as:

$$TieStrength(u_i, u_j, \triangle t) = \frac{\#(ForeignInteractions(u_j, \triangle t) \cap DirectMentionInteractions(u_i, \triangle t))}{\#(DirectMentionInteractions(u_i, \triangle t))} \quad (15)$$

$u_i \in U, u_j \in SN(u_i)$ Basically, tie strength from user u_i on user u_j is the ratio between the interactions created by u_i where u_j has been particularly mentioned and all interactions created by u_i mentioning somebody. $TieStrength(u_i, u_j, \triangle t)$ is not necessarily $TieStrength(u_j, u_i, \triangle t)$. This metric is supported by the definitions Def. 6 and Def. 7.

Instead of taking the subset $DirectMentionInteractions(u_i, \triangle t)$ in the previous definition, it would be also possible taking the entire set $Interactions(u_i, \triangle t)$, but the tie strength would return rather lower numbers, as many users just broadcast messages to their entire network without explicitly mentioning anybody in particular. The Figure 1 shows a fictive time line over 4 days for the users X,Y and Z and 3 other users A,B and C in $SN(X)$. The values for the metrics required to compute the Tie Strength for this example can be found in Figure 2 with the entire set of combinations for the SM users A,B,C,X,Y and Z.

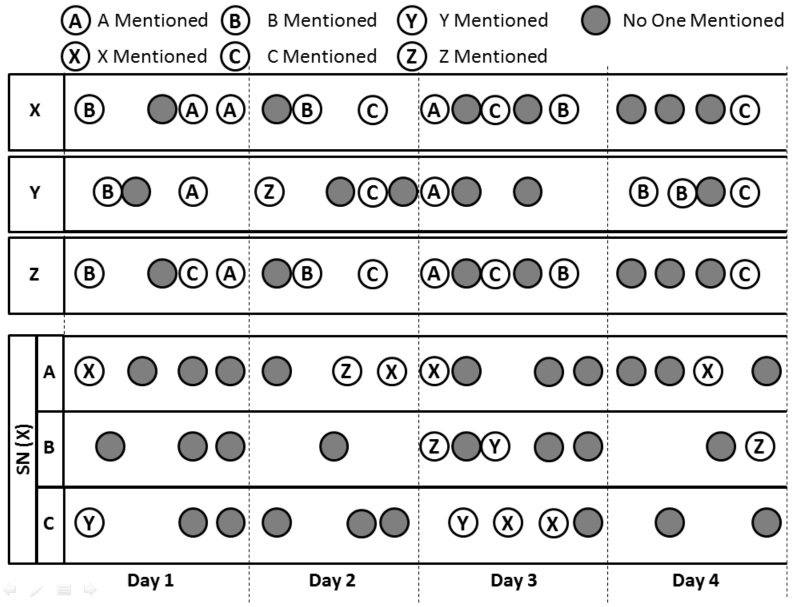

Figure 1: Social Media Tie Strength Computing Sample

	Interactions			TieStrength					
	Total	Direct Mentions	Foreign	X	Y	Z	A	B	C
X	16	9	6	-	0,00	0,00	0,33	0,33	0,33
Y	14	8	2	0,00	-	0,13	0,25	0,38	0,25
Z	16	9	3	0,00	0,00	-	0,22	0,33	0,44
A	15	5	7	0,80	0,00	0,20	-	0,00	0,00
B	11	3	8	0,00	0,33	0,67	0,00	-	0,00
C	12	4	9	0,50	0,50	0,00	0,00	0,00	-

Figure 2: Tie Strength Metrics based on the example in Figure 1

3.5. Intrinsic and Extrinsic Impact Metrics

Our suggestion for modelling the impact created by an particular user in a place over his/her SM channels relies on 2 components: the first one focuses on just the user's behavioural aspects and posted SM content —intrinsic component— whereas the second one takes into account the interaction with the SM network of the user —extrinsic component—.

Based on the *Differential Perception Factor* and the *Entity Engagement Index*, we define the intrinsic component:

$$Intrinsic\ Impact(u_i, E, P, \triangle t) = \Im(EEI(u_i, E, \triangle t), DPF(u_i, E, P, \triangle t)) \tag{16}$$

The extrinsic component requires the joint computation of the *Exposed Users* set and the *Tie Strength*:

$$Extrinsic\ Impact(u_i, E, P, \triangle t) = \Im(ExposedUsers(u_i, E, \triangle t), TieStrength(u_i, SN(u_i)) \tag{17}$$

Which can be implemented as an addition of the Tie Strength with u_i of all users in the exposed to the interactions created by u_i in the time period under analysis:

$$Extrinsic\ Impact(u_i, E, P, \triangle t) = \sum_{j=0}^{\#ExposedUsers(u_i,E,\triangle t)} TieStrength(u_i, u_j, \triangle t) \tag{18}$$

The resulting impact is then a combination of both intrinsic and extrinsic components:

$$Impact(u_i, E, P, \triangle t) = \Im(IntrinsicImpact(u_i, E, P, \triangle t), ExtrinsicImpact(u_i, E, P, \triangle t)) \tag{19}$$

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

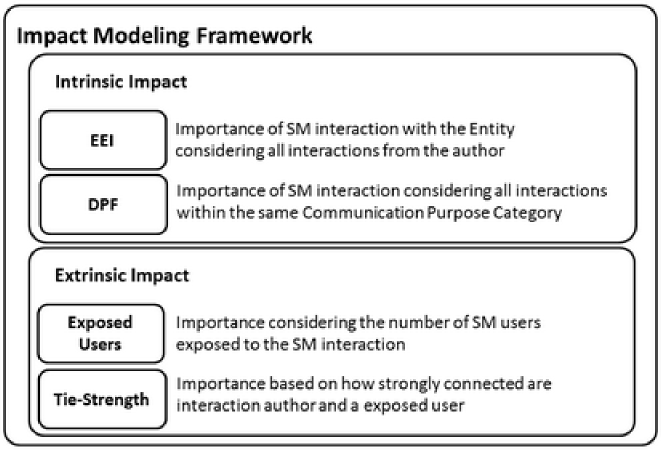

Figure 3: Overview of the meaning of the metrics defined in our framework

3.6. Making the results actionable

The underlying complexity to the metrics computing might compromise the overall performance of the system, delivering highly accurate results but not quick enough to take decisions upon. Thus, we provide ways of obtaining actionable results in shorter time when the use case forces the trade-off between accuracy and time-to-results to be decided in favour of the later. Higher precision implies higher latency, which might be appropriate for batch analysis, but not meet the requirements for an early-warning fast-reaction system. The simplifications introduced in this section are designed in a way that the metrics' values inflate —increase of false positive situations—, which from the business perspective is more acceptable than the other way around (rather alerting on something that maybe is not that important than not alerting about an important situation at all).

We have identified the complexity drivers and suggested alternative ways of computing the previously defined metrics (see Figure 3) when the time to results is more critical. Our approach to approximation for the before presented metrics is described below:

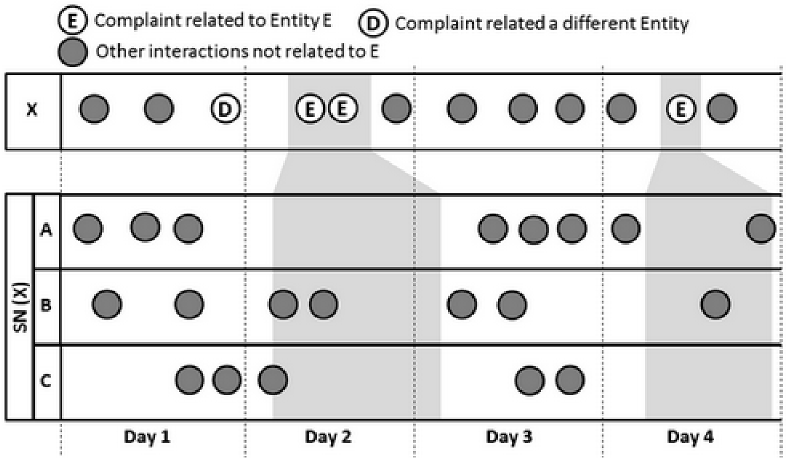

Figure 4: Fictive Social Media Interactions Set to illustrate the Impact computation

			A	B	C
DPF (X, E, P, Δt)	0,75	***Tie Strength (X,***	0,10	0,50	0,20
	day 1	**day 2**	**day 3**	**day 4**	**Aggregated /Averaged**
ExposedUsers(X, E,	Ø	{B,C}	Ø	{A,B}	{A,B,C}
# ExposedUsers(X, E,	0	2	0	2	3
EEI(X, E,	0	0,67	0	0,33	0,25
IntrinsicImpact (X, E, P, Δt)					
Avg	0	0,35	0,00	0,30	0,1625
Sum	0	0,70	0,00	0,60	0,325
ExtrinsicImpact (X, E, P,Δt)	0	0,50	0,00	0,25	0,1875
Impact(X,E,P,Δt)					
Avg	0	0,18	0,00	0,08	0,175
Sum	0	0,35	0,00	0,15	0,125

Figure 5: Impact Metrics based on the example in Figure 4

3.6.1. Intrinsic Impact approximation

Removing the SM behavioral bias is *DPF's* main job, but computing it requires fairly complex time-consuming NLP operations to assign each and every SM interaction made by the SM user to the appropriate communication purpose. The *DPF* can be simplified as follows at the expend of keeping the potential SM behavioral bias:

$$DPF(u_i, E, P, \triangle t) \approx 1 \tag{20}$$

Alternatively, the system could store and return a counter for each communication purpose category for the user, under the assumption that the SM behavioral bias does not strongly change over time. It would substantially simplify the computation of the *DPF* metric.

The *EEI* requires the scanning of the latest interactions created by the user u_i and the flagging of those that are related to the Entity *E*. This step can be spared by approximating the EEI by a value specific to the user, to the location or just generic for all users. As the engagement with an Entity *E* can strongly vary driven by events of all kinds, the use of a pre-computed *EEI* value for a given user u_i based on historic data might lead to slightly less accurate results.

3.6.2. Extrinsic Impact approximation

Unlike the components of the Intrinsic Impact, Exposed Users and Tie-Strength require a joint simplification, as both refer to the user authoring the SM interaction u_i and the user being exposed to it $u_j, u_j \in SN(u_i)$. Computing the set of exposed users requires extracting all interactions of the entire $SN(u_i)$ during the period $\triangle t$ for further computing of the exposure window for each user in $SN(u_i)$. The Tie-Strength calculation is performed by analyzing all direct mentioned interactions and the foreign interactions for the involved pair of users, which again requires the pulling and scanning of all interactions during the period under analysis.

Both Tie-Strength and Exposure are adjusting factors of the total number of followers the user u_i has in his/her Social Network. To avoid the single computation at user level of Exposure window and Tie-Strength, we can work with predefined value distributions, dividing the Tie-Strength values range into intervals and multiplying the total number of followers of u_i by the proportion our distribution function assigns to each interval. These distributions can be based on frequency of occurrence by value interval.

$$TieStrength(u_i, \triangle t) \approx \#SN(u_i) * \sum_{k=0}^{\#D} k * d(k),\ d\ defined\ by\ D \tag{21}$$

Where d is the chosen distribution consisting of $\#D$ intervals, $d(k)$ is the value associated with the interval k and $\#SN(u_i)$ is the cardinal of the followers of u_i. For example in Figure 6, the distribution D is defined in 4 intervals $k \in 5, 10, 22, 63$ with following weights $d(5) = 1, d(10) = 0.75, d(22) = 0.5$ and $d(63) = 0.25$. In the example, $\#SN(u_i) = 135$, which results in a TieStrength of $135 * 0,2143 = 28,9305$. The TieStrength value is always bigger or than 1; if for particular privacy settings our crawler cannot retrieve the number of followers or just because a particular user does not have any follower, we assume that every user at least follows him/herself.

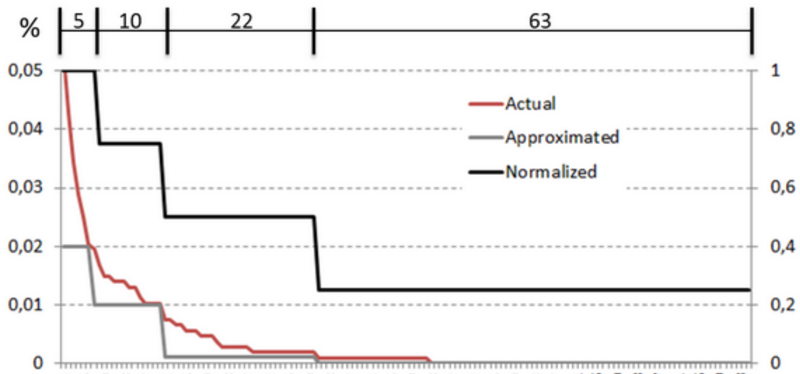

Figure 6: Example of Tie-Strength approximation by a weighted distribution

The same procedure can be applied for approximating the cardinal of the ExposedUsers set:

$$\#ExposedUsers(u_i, \triangle t) \approx \#SN(u_i) * \sum_{k=0}^{\#D'} k * d'(k),\ d'\ defined\ by\ D' \tag{22}$$

Where d' is the chosen distribution consisting of $\#D'$ intervals, $d'(k)$ is the value associated with the interval k and $\#SN(u_i)$ is the cardinal of the followers of u_i.

This approach still reflects the differential link strength distribution within a given user's social network, but does not differentiate users overly tied to their networks from users almost not mentioning anybody in their tweets (just broadcasting information without ever meaning anybody). It could be achieved by introducing a weighting factor in the distribution $d(k) * w(k)$, however, computing this weighting factor requires analysing case by case the Tie-Strength of the users u_i with each user $u_j, u_j \in SN(u_i)$. In this case, the complexity involved is similar to the one required for computing the original Tie-Strength metric. This can be achieved by account ranking or other type of graph SM ranking.

Another alternative for approaching $\#ExposedUsers$ is just multiplying the $\#SN(u_i)$ by a factor $\#ExposedUsers(u_i, \triangle t) \approx SN(u_i) * K,\ K \in [0, 1]$, which would be consistent with the binary character we expressed in the definition of Exposed Users (subsection 3.4). This option is not the best choice for approximating the Tie-Strength, as it would no longer reflect how differentially strong the links between different users are.

3.6.3. Working with Levels

The complexity of decision making for business stakeholders based on a priori resulting large numbers quantifying the impact might make the CARESOME output difficult to consume. Thus, we suggest the mapping of the impact metric values to categories, as many as different action making scenarios are defined in the use case or are meaningful for the business. Each category or level is defined by a *min* and a *max* value for the metrics. Typical category schemes are the traffic light inspired RAG (Red, Amber, Green), some Likert-inspired [60](e.g.: *Very Strong, Strong, Medium-light, Light, No Impact*).

The suggested *Levels* might no provide similar results if defined globally. Sometimes, setting up the defining Levels max-min pairs specifically for a location soften the differences between geographical areas.

4. CARESOME System Architecture

CARESOME is designed to pull the SM content generated for a set of predefined locations over time and measure the impact of all SM interaction on a defined Entity (company, institution, brand, etc). These metrics can then be used to understand when a customer retention campaign is required in a particular location, when the entity's image is damaged or weakened in the location and competitors can execute promotional actions with higher conversion chances, etc. Additionally, CARESOME offers (after proper configuration), the possibility of monitoring similar impact metrics on immediate competitors, which allows for identifying weak points on locations that can be targeted more aggressively by acquisition campaigns.

Figure 7 shows the modules of the system as well as their input data sources and the data storage used for the information exchange between them (*Tweets Harvester, Tweets Classifier, User Data Collector, Metrics Generator*). In the following subsections we are going to describe each and every step from the data gathering to the metrics presentation stage, explaining which modules are involved and providing details about the implementation.

4.1. Tweets Harvester

Relying on the Twitter Search API [1], the component *Tweets Collector* periodically extracts all SM interactions generated in a SM location and stores them into a database for further processing. The location is typically defined as a pair of geo-location coordinates —latitude and longitude— and a radius. A pre-filtering by language can also be applied to the harvester to just pick tweets in a given language. The Geo-Gazetteer and Geo-Coder components help allocation SM interaction with missing GPS coordinates to the right areas.

4.2. Tweets Classifier

The role of the classifier consists of flagging the previously gathered tweets that are related to the Entity we are analyzing on one hand, and assigning a Communication Purpose Category to these tweets on the other hand. The *Entity Flagger* component is configured by the so called *Entity Definition File*. All potential terms pointing out the relationship between the tweet content and the Entity shall be part of this file. These terms can be names of SM accounts —like the official brand account, the entity news account, the entity accounts specific to a country, etc—. For example, taking the airline *British Airways* from UK, we would have the Official British Airways Global account @British˙Airways, the account for the North American Customers Care @BristishAirways, the accounts related to official and unofficial news and press releases related to the airline @BA_Headlines and @BritishAirNews, accounts for "haters" like @We_hate_BA, etc. All relevant hash-tags shall also be included, for example the ones used by

[1] Available at `https://dev.twitter.com/docs/api/1/get/search`

the company for running marketing campaigns (*#UnGroundedThinking*,etc), the ones referencing the entity itself (*#BAirlines, #BritishAirlines, etc*, the ones defined by customers to spread their lack of satisfaction (*#BAsucks*), etc. Also the name of the services offered by the company and /or name of the products —in our example, flight numbers like BA0177, etc—. Depending on the scope of the analysis, sub-brands might be also part of this file (such as @flybmi for the British Midland International airline). Additionally, typical n-grams with for example the slogan of the company or of a particular campaign, etc are included (e.g.: "Learning to fly").

As in SM due to the brevity but also due to the typing speed, the spelling mistakes are quite frequent —getting even worse with the adoption of small screen devices and the sometimes unwanted effect of the automatic spelling corrector—, our flagging component implements a tolerance threshold given by a string similarity function [61] to accept spelling mistakes (e.g.: *birtish airways* or *british airways* with a similarity over 0.7 wouldn't be rejected if the threshold was set to 0.6)

The *Communication Purpose Flagger* works according to a similar input source (a definition file containing the terms for identifying a communication purpose category), but applies a more complex process. Each geo-located tweet is tokenized applying a sentence tokenizer first and a word tokenizer later (based on [62]) both adapting the Punkt Tokenizer [63] to deal with social media texts.The modified tokenizer provides the stop words removal as well, so that a lemmatizer takes over. The lemmatizer extracts the lemmas we then match against the input definition file. Both number and definition of categories depend on the particular business needs. For example, if there is a department specialized in handling complaints, a separate one running retention and acquisition campaigns, a third one in charge of improving the brand index, etc. makes sense understanding which communication purpose category maps to which action plan to be taken by which department and define them accordingly. In situations where a simple monitoring does the job, a sentiment-based separation of the tweets is sufficient.

Finding the defining terms for a particular communication purpose category is challenging because of the potential overlapping with another category and because of the underlying complexity in the Natural Language Processing. Analysing previously generated user content related to specifically a given category in channels like an online Forum, a product review section, etc and performing n-gram extraction tasks [64, 65] over a long history can help identifying the defining terms.

The *Classifier* module also implements a simple disambiguation mechanism relying on both Part of Speech tagging and the presence of more than one terms related to the Entity or Purpose Category. Additionally, for especial cases where the ambiguity impacts the name of the Entity E, a Naive-Bayesian classifier sufficiently trained helps separate senses, so that only tweets related to the Entity are flagged (e.g.: *Emirates* can be the name of the airline, but also the name of the Arsenal Stadium or even the country —UAE—).

4.3. User Data Collector

The purpose of this module is extracting all the information related to the authors of the flagged tweets to enable the computing of all relevant metrics. It's divided into different components, each one addressing a particular data gathering task. Each component can be also configured to apply just the data gathering required for the approximation described in the Subsection 3.6.3, instead of a more thorough yet slower data gathering required for the full-fledged metrics. The User Data Gathering Components are described below.

4.3.1. User Data Retriever

All relevant information about the user provided by the Twitter API, including number of followers, friends, retweets, etc are retrieved and persisted by this component. The system supports the filtering of certain accounts with a black-list mechanism, meaningful to exclude for example the company's employees accounts or the company's SM department accounts.

4.3.2. Network Data Retriever

To compute the Extrinsic Impact metrics, all kinds of information related to the Social Network of the author of any of the SM interactions related to the Entity of interest are required. This component extracts the entire SN framework for the identified SM authors in the previous module and persists them with a time-stamp. It allows for handling changes over time (new followers, followers leaving, etc) which becomes especially critical when for performance reasons, the Tie-Strength is computed once per user and used going forward.

This component can also be configured to not retrieve anything if the both Tie-Strength and ExposedUsers are going to be approximated as explained in the subsection 3.6.3. The number of followers, which is the only input that is really required is made available by the previous component, as explained above.

4.3.3. Time-line Extractor

Extracts the latest X tweets created by the Author, as well as the X latest tweets created by each user in the SN of the author to later enable the computing of Tie-Strength, as well as the exposure window.

4.4. Metrics Generator

Once the set of required data has been gathered, processed and stored, the impact metrics are computed and transformed to be consumable in decision making scenarios. This is the purpose of this module, which relies on following three components.

4.4.1. Intrinsic Impact Calculator

The Entity Engagement Index and the Differential Perception Factor for the authors of the SM interactions related to the Entity under analysis flagged by the Tweets Classifier are created by this component. For all authors of the flagged tweets, the *EEI* is computed applying the formula 10. Similarly, the *DPF* is obtained as per the formula 12 for later combination of the results, as defined by the Intrinsic Impact equation 16.

If the system is set up to apply the approximations defined in Section 3.6, the DPF doesn't need to be computed (as 1 or other number close to 1 is always taken). Likewise, the EEI is taken as a configured value, which can be a generic one valid for all users or specific to a location which has been previously entered in the system and held in a look-up table (e.g.: the same for cities with a population between 50K and 100K). Alternatively, this module can implement a look-up table where the EEI is kept at user level from previous runs. The system can be set up in a hybrid mode, so that no approximation is done for users not present in the look-up table, but the values existing in the table are used as a kind of caching mechanism.

4.4.2. Extrinsic Metric Calculator

For the SM users who authored the tweets flagged as related to the Entity, both Tie-Strength and Number of Exposed Users are computed in this module. Applying the formula 15 over a set of current interactions created by the user u_i, the Tie-Strength with each and every member of $SN(u_i)$. The number of interactions considered in this set can be fixed (e.g.: the latest 1000) or can be dependent on a timely factor (e.g.: all interaction in the last 3 months). The larger the set of interactions, the more accurate the Tie-Strength but also the higher the risk of neglecting decaying Tie-Strengths (user that used to have a very close interaction-rich relationship in the past but no longer at present). Similarly and as explained in the 3.4, the set of exposed users is computed with the formula 14

If CARESOME is configured in speed modus, the approximations explained in the Subsection 3.6.2. Depending on the methods configured for the approximation, CARESOME applies the weighted distribution explained in the equation 21 for Tie-Strength and in the equation 22 for the Exposure. Both length and weight need to be configured in the system. If the decision is to keep *memory* on previously calculated Tie-Strength values between users, the system provides the lookup table to support this process.

Once the individual Impact has been computed, the Metrics Generator module computes the overall Impact the aggregation as explaining in the formula 19).

4.4.3. Output Category Mapper

To make the impact values per communication purpose category actionable, CARESOME provides a dynamic mapping of values to categories, whose min and max values automatically adjust based on the values distribution. The number of categories is configurable, as well as the time granularity the impact value is provided for (as explained in the Subsection 3.5). In the next section, we show concrete examples of this mapping applied (see also Figures 17, 15)

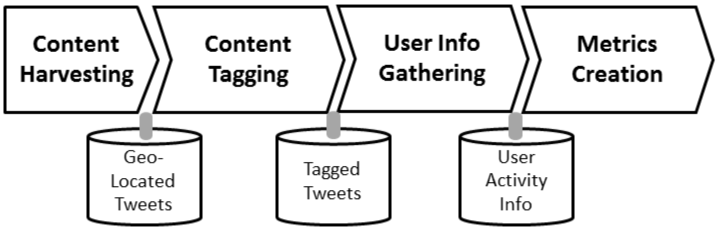

Figure 7: System structure

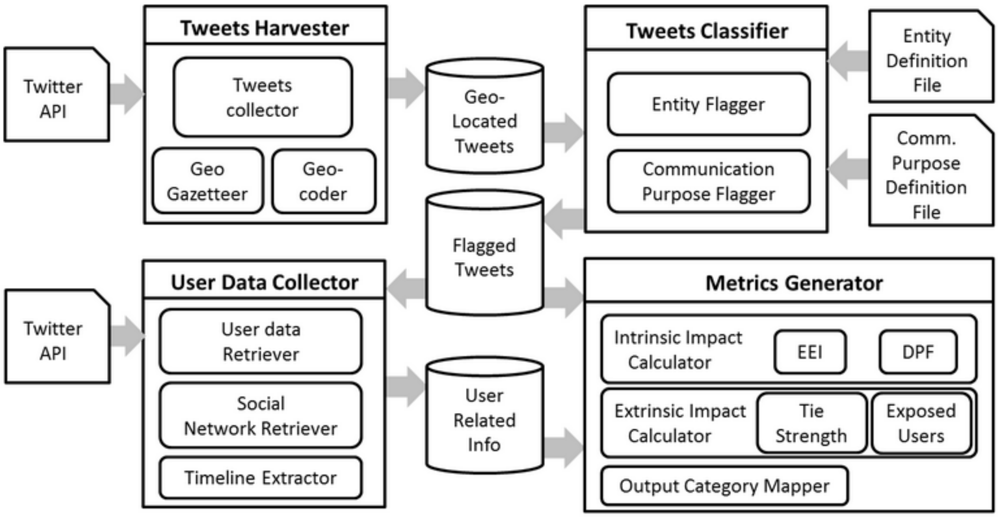

Figure 8: CARESOME System Modular Architecture

5. Analysis of performance and discussion

To analyse the performance of our system in action we chose 2 well-known locations with a high volume of visitors and where people are likely to have time and therefore prone to create Social Media interactions: the biggest two airports in the city of London, namely Gatwick and Heathrow. We set 2 harvesters centred in the middle point of both airports with a radius big enough (5 km) to capture all activity happening at both airports (see Figure 9).

To increase the number of interactions retrieved by the geo-location query, we also ran 2 harvesters configured to just gather tweets with the words *Gatwick* or *Heathrow* present. Thus, we were able to gather an additional set of interactions from those users that didn't have the geo-location functionality enable but referred to one of these airports.

Between the 24ˆth of November 2013 and 23ˆrd of January 2014, 852319 SM interactions have been gathered. During this period of time there were severe weather conditions, spreading the chaos all over the country with strong winds and flooding episodes, which impacted the quality of all transportation services in UK. Thousands of passengers were affected and the Social Media platforms filled with users' statements on how well the different carriers handled the incident.

For our show case we took as entities a subset of the airlines operating in these airports and gathered the identifying terms (see Figure 10).

As Communication Purpose Categories, we worked with the standard ones: *Complaints and Criticism (c)*, *Praise and Positive Feedback (p)*, *Information Request and Customer Care (ir)* and a forth one for the rest called *Neutral (n)*. The creation of the definition files for these categories has been performed by enhancing a pre-defined default file with typical terms for each category with the most frequent terms in manually flagged airline specific Tweets. Figure 11 shows the terms per category sorted by frequency over all airlines.

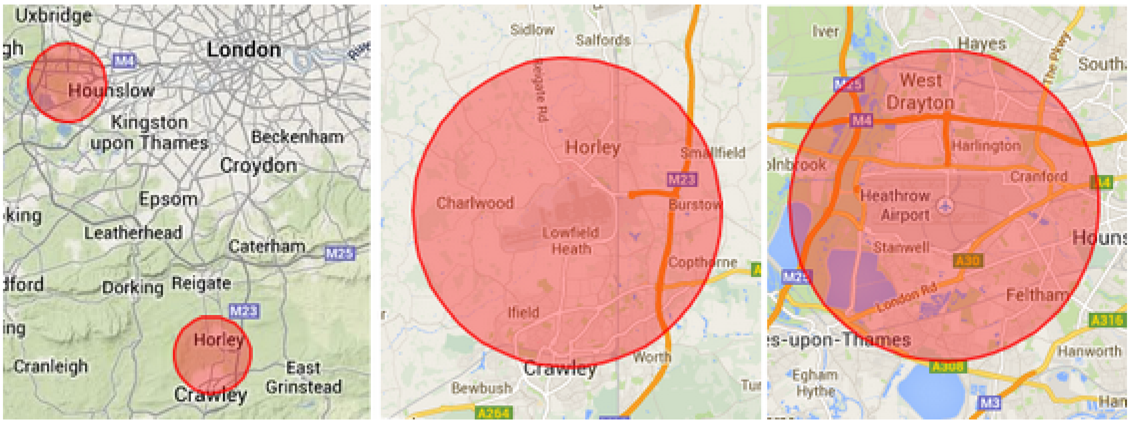

Figure 9: Harvesters overview

aegean	**air india**	**british airways**	**etihad airways**	**Kenya Airways**	**qatar**	**Thomson Airways**
aegean	air india	british airways	etihad airways	Kenya Airways	qatar	Thomson Airways
aegeanairlines	**air lingus**	British_Airways	EtihadAirways	KenyaAirways	qatarairways	**transavia**
aeroflot	AerLingus	BritishAirways	**finnair**	**klm**	**Ryanair**	transavia
Aeroflot	lingus	**brussels airlines**	finnair	klm	Ryanair	**tunisair**
aeroflot	**Air New Zealand**	brussels airline:	**flybe**	KLM_EIR	**SAS**	tunisair
aeromexico	Air New Zealand	FlyingBrussels	Flybe	KLM_SE	SAS	**Turkish Airlines**
aeromexico	FlyAirNZ	**cathay pacific**	**germanwings**	KLM_UK	**singapore**	turkish
air asia	**alaska air**	cathay pacific	german wings	KLM_US	singapore	TurkishAirlines
AirAsia	alaska air	cathaypacific	germanwings	klmfan	Singapore Air	**united**
air baltic	AlaskaAir	cathaypacificUS	**Hawaiian Air**	**korean air**	SingaporeAir	united airlines
airBaltic	**alitalia**	**czech airlines**	Hawaiian Air	Korean Air	**south african**	UnitedAirlines
air berlin	Alitalia	Czech Airlines	HawaiianAir	KoreanAir	south african	**us airways**
air berlin	**american air**	CzechAirlines	**Iberia**	KoreanAir_KE	**swiss airlines**	us airways
airberlin	American Air	**delta airlines**	Iberia	**lufthansa**	FlySWISS	USAirways
airberlin_com	AmericanAir	Delta	Iberia_en	lufthansa	**TAM Airlines**	**virgin**
airberlin_US	**american airlines**	DeltaBlog	**Icelandair**	Lufthansa_DE	tam airlines	virgin
air canada	american airline:	DeltaNewsroom	Icelandair	Lufthansa_USA	TAMAirlines	VirginAmerica
Air Canada	**asiana**	DeltaSkyBonus	**japan airlines**	**Midwest Airlines**	**tap**	VirginAtlantic
AirCanada	asiana	DeltaTechOps	japan airlines	Midwest Airlines	tap airlines	VirginAustralia
air europa	AsianaAirlines	**easyjet**	**jetairways**	MidwestAirlines	**Thai Airways**	**vueling**
Air Europa	Flyasiana	easyjet	jetairways	**monarch**	Thai Airways	vueling
AirEuropa	**atlantic airways**	easyjetservice	**JetBlue**	Monarch	ThaiAirwaysAust	
air france	Atlantic Airways	**emirates**	JetBlue	**norwegian**	ThaiAirwaysIT	
air france	AtlanticAir	emirates	**Jetstar**	Fly_Norwegian	ThaiAviation	
AirFranceFR	AtlanticJet	**estonian air**	Jetstar_Japan	Norwegian	**Thomas Cook**	
AirFranceIE	**austrian airlines**	estonian air	JetstarAirways	**qantas**	Thomas Cook	
AirFranceUK	austrian air	Estonian_Air		qantas	ThomasCookUK	
AirFranceUS				QantasAirways		
				QantasUSA		

Figure 10: SM accounts and hashed tags used to identify the major Gatwick / Hethrow airlines

For Tie-Strength and Exposure, CARESOME was configured to rely on the weighted distribution presented in the Figure 6.

Figure 12 shows the result of the classifier per airline and per harvester. As expected, airlines just operating from one airport present a much higher amount of interactions in this airport (e.g.: *Easyjet* showing just a few interactions in Heathrow compared with Gatwick). An interesting exception is *British Airways*, as it is used as reference in opposition to low-cost carriers in Gatwick, even if no BA flights departures from or lands there.

The adverse weather conditions on the 24ˆth and the 25ˆth of December left thousands of passengers stranded in

complaint			praise			info request		
term	#	share	term	#	share	term	#	share
delay	74	7,64%	great	44	14,29%	can you	35	50,72%
cancel	63	6,51%	good	38	12,34%	which	19	27,54%
no info	52	5,37%	love	29	9,42%	where	4	5,80%
late	39	4,03%	thank you	29	9,42%	do i	3	4,35%
bad	34	3,51%	nice	22	7,14%	do you	3	4,35%
miss	29	3,00%	best	21	6,82%	how do	2	2,90%
stranded	29	3,00%	lovely	16	5,19%	can i	1	1,45%
ruin	28	2,89%	impressed	15	4,87%	can u	1	1,45%
stuck	28	2,89%	better	13	4,22%	do u	1	1,45%
poor	27	2,79%	amazing	12	3,90%			
chaos	24	2,48%	awesome	12	3,90%			
fail	23	2,38%	well done	11	3,57%			
lose	18	1,86%	excellent	9	2,92%			
break	17	1,76%	cool	7	2,27%			
no staff	17	1,76%	lucky	6	1,95%			
worst	17	1,76%	favourite	4	1,30%			
problem	16	1,65%	loving	4	1,30%			
other	433	44,73%	other	16	5,19%			

Figure 11: Top 30 terms for purposes identification

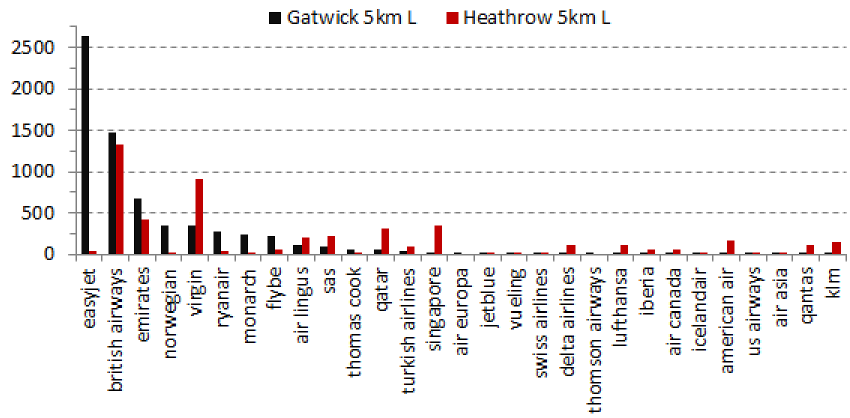

Figure 12: Number of flagged interactions per airline and location for the top 30 airlines

the Gatwick Airport due to power problems [2]. Countless flights were canceled or suffered severe delays [3]. In this emergency situation, a blame game between Gatwick airport and the airline Easyjet started [4]. In Figure 13, we can see the daily values for the single impact metric components, both intrinsic and extrinsic for the entity *Easyjet* and the communication purpose category *Complaints* measured by the Gatwick harvester. Especially on the 24th we observe a peak over all sub-metrics, motivated by the increase of SM interactions (297 different users) criticizing the way *Easyjet* handled the emergency situation. These results produced by CARESOME would have given *Easyjet* enough quantified evidence to trigger some sort of reaction and the corresponding communication back to the SM channel to palliate the incident effect. After the potential airline reaction, CARSOME can then measure the SM community response. In general, CARESOME's role is providing enough insights for a company to steer the SM dialogue in all fronts.

On the Christmas eve, due to the disruptions in the railway transportation, many passengers were about to miss

[2] `http://www.bbc.com/news/uk-england-sussex-25503513`

[3] `http://www.itv.com/news/story/2013-12-25/gatwick-airport-christmas-travel-disruption-cancellations/`

[4] `http://www.dailymail.co.uk/travel/article-2535822/Blame-game-Gatwick-easyJet-clash-responsibility-Christmas-Eve-chaos.html`

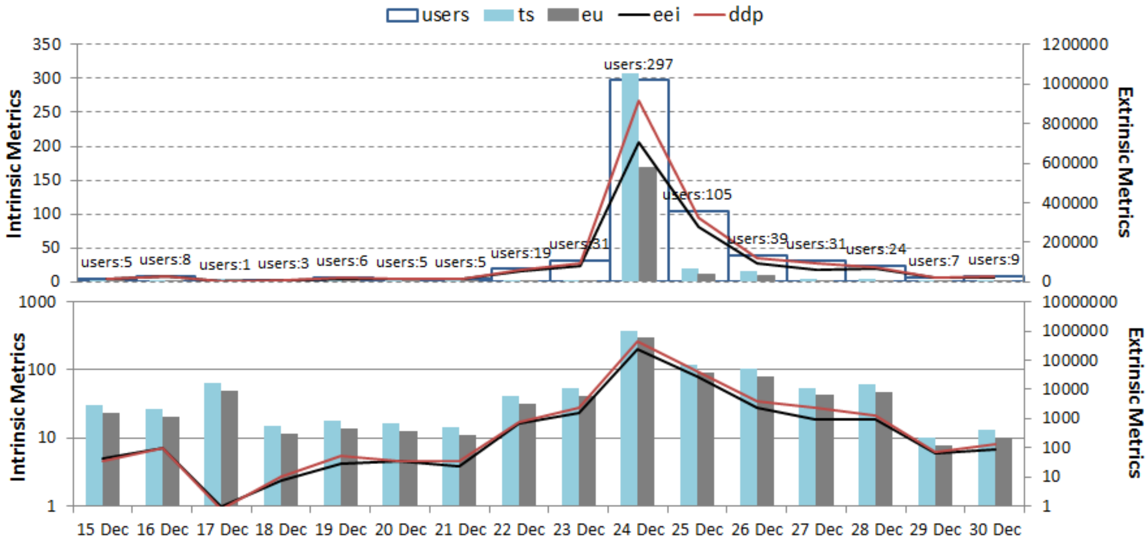

Figure 13: Metrics components (TS, EU, EEI and DDP) as defined in Section 3.5 for Easyjet in Gatwick over 2 weeks. Natural (above) and logaritmic (below) scale

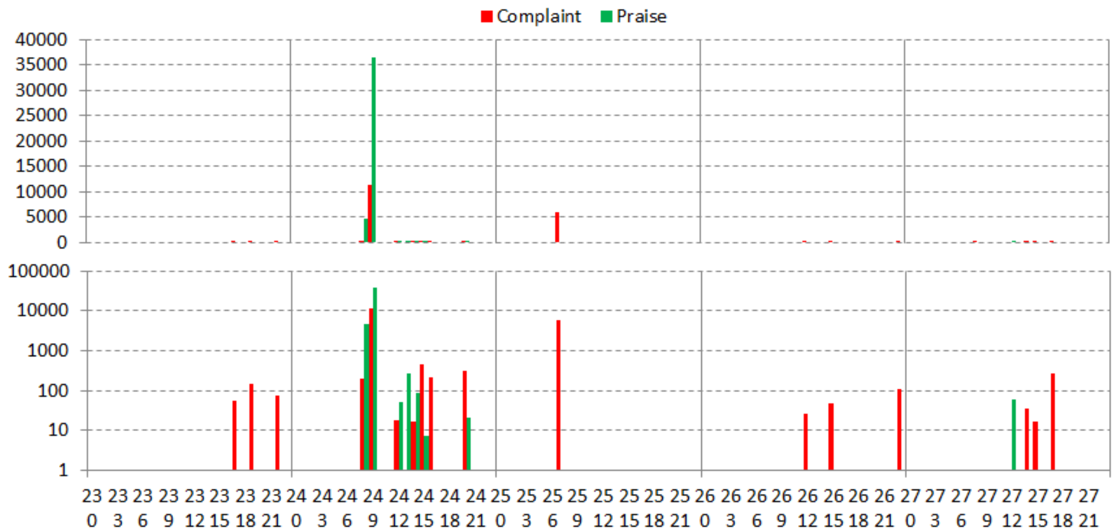

Figure 14: Hourly impact for Ryanair in Gatwick between the 23ˆrd and 27ˆth Dec 2013

their flights departing from Gatwick. *[...]Ryanair uses the South Terminal, but decided to delay its services by an hour to Cork, Shannon and Dublin by an hour "To ensure all those affected by rail delays at Gatwick get home [...]".* [5]. CARESOME reported a peak in the impact created by SM interactions (more than 85 in total) talking about it in the communication *Praise* purpose, with messages such as: *"I must say @Ryanair handled everything really well yesterday at gatwick #gladtobehome"*. Figure 14 shows the peak at 9:00 am the 24ˆth, consequence of all praise-related interactions. Ryanair delayed its flights to allow people get home for the Christmas eve and such a small decision had a huge impact over the SM channels as reported by our system, outperforming even the bad press related

[5] `http://www.independent.co.uk/travel/news-and-advice/passengers-stranded-at-gatwick-airport-as-flooding-causes-power-outages-9023990.html`

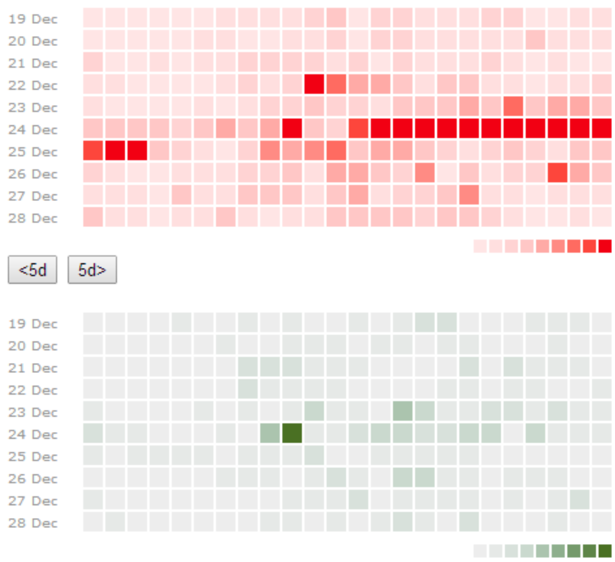

Figure 15: 10-days hourly heatmap for all airlines for categories *Complaint* and *Praise* for the Gatwick Harvester

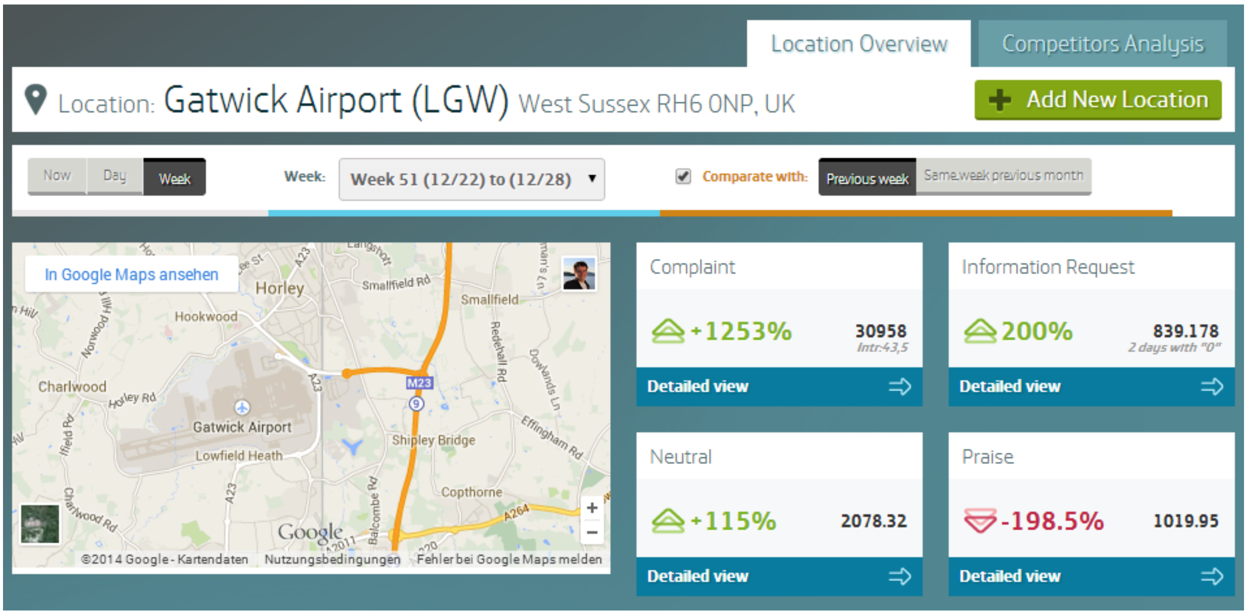

Figure 16: Weekly dashboard view for *Easyjet* in the Gatwick location

to service disruption. This example shows very well when finer time granularity (hourly instead of daily) makes sense and how the impact measured by CARESOME delivers meaningful results aligned with one event that happened in the real world and got reflected in the SM channels.

CARESOME can also help understanding and measuring those small things that might not be considered by the company as relevant for its customers but perceived by those as such. A captain successfully landing a plane after complicated maneuvering with adverse weather conditions might be seen as part of his job, but might also trigger a set wave of SM interactions praising the action [6] (which also contributed to the increase in the category *Praise* on

[6] Video recorded by a passenger showing the captains' heroic landing `https://www.youtube.com/watch?v=MPT3bdEr_VM`

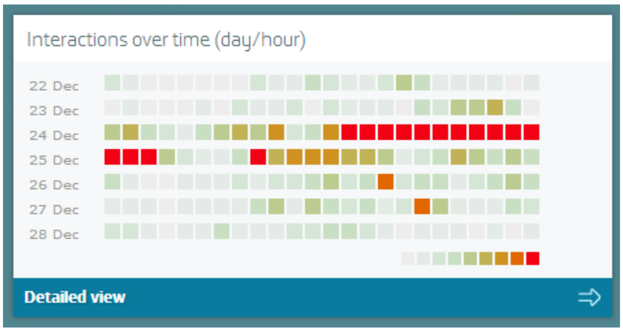

Figure 17: Weekly Social Media Interactions Overview for Gatwick Airport

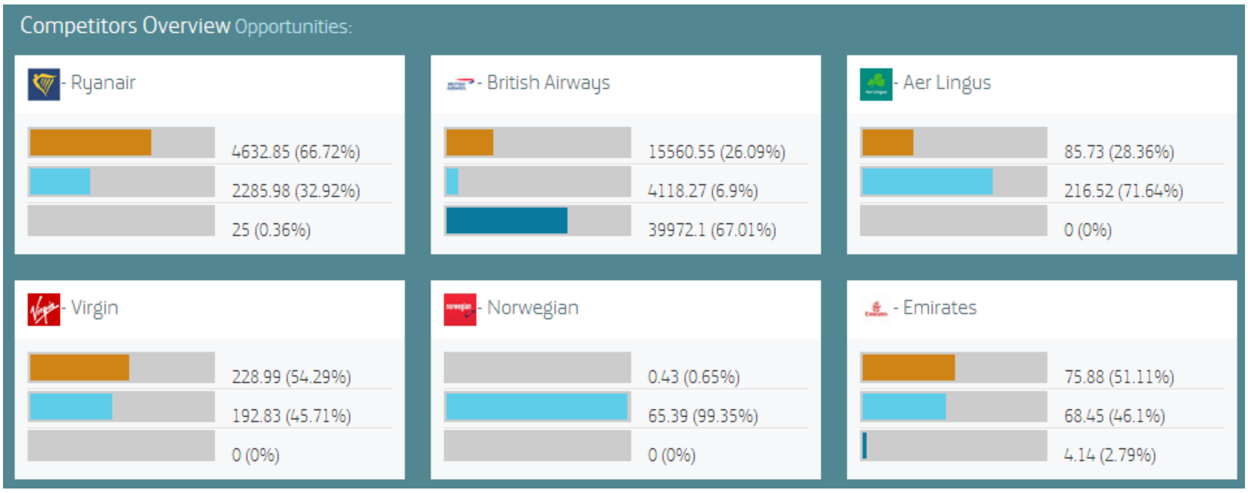

Figure 18: Weekly Easyjet competitors' overview for Gatwick Airport: C (orange), P (light blue), IR (dark blue)

the 24ˆth Dec. in Heathrow for *British Airways* as we can see in Fig. 19). As a potential take to action, British Airways might have well created and launched a campaign to reinforce the idea of security in extreme conditions. With CARESOME, the impact of this campaign could also be measured or any other public relationship action in an ongoing basis.

Fig. 16 shows the system cockpit for *Easyjet* in Gatwick with the time granularity set to one week (in this case the week starting 22nd Dec)[7]. In addition to the impact values for each category, the change from the previous week in percentage helps understanding whether something exceptionally changed which requires some kind of reaction by the brand. As our impact metric can deliver pretty high numbers, a heatmap-like visualization over time units allows for a quick visual identification of high-impact increases (see Fig. 15 displaying hourly values for complaints and praise over 10 days). Clicking on a particular square for a given day and a given hour displays the interactions flagged for the communication purpose category that took place there. To better makes sense of the impact values, the dashboard offers as well a calendar heat-map showing the number of interactions per hour (see Fig. 17). Apart from tailoring retention campaigns on locations where for example the impact of complaints substantially increases or keeps increasing, CARESOME can be used to monitor the perception of direct competitors in a given location to spot acquisition opportunities. Figure 18 is a snapshot from the system front-end showing the category share per competitor for *Easyjet* in Gatwick over the 51th week of the year.

[7]Pictures for logos and airline data have been taken directly from Twitter for the purpose of this research

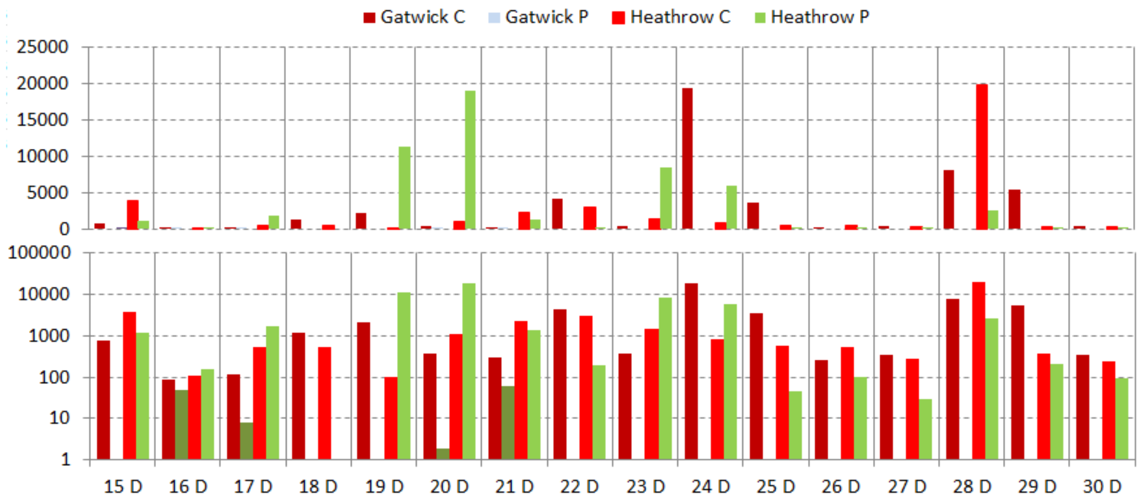

Figure 19: 15 days monitoring of British Airways impact for Complaints and Praise in Gatwick and Heathrow - Natural and Logaritmic scale

5.1. Extreme cases analysis

Analyzing how the metrics perform in extreme cases helps understanding both sensibility and suitability for real-world business scenarios. Let's assume following setup for a given user u_i over a period of time $\triangle t$:

- *All parameters maximized*: which means:
  - *DPF* is (close to) 1: the only complaint the SM user posted was about the Entity and otherwise, the user does not post any complaint.
  - The *EEI* is also 1: the user just posts about the Entity and nothing else.
  - The *Tie Strength* between the user u_i and any user in $SN(u_i)$ is also 1: all Foreign Interactions of any of the users in $SN(u_i)$ have been done by user u_i. In other words, u_i got the full attention of any user in $SN(u_i)$.
  - All users in $SN(u_i)$ has been exposed to the interactions of u_i.

  In this case, the value of or Impact metric is equal to the size of the Social Network of the user u_i. $Impact(u_i, E, P, \triangle t) = \#SN(u_i)$
- *Low User Entity Engagement*: overly active users posting a lot of content to their SM Networks and just engaging once with an Entity E to express a complaint, ask a question, etc. are penalized over rather passive users who turn active to engage with E.
- *Overly complaining users*: when most of the interactions of user u_i belong the the same Communication Purpose Category (e.g.: always complaining, always expressing a "kudos" or a "well done", his/her network tends to lower the perceived impact of the interaction. The Differential Perception Factor helps modulating the impact metric based on the user's SM behavior; the impact of an overly complaining user posting a complaint about an Entity is lowered down according to the DPF.
- *User loosely tied to his/her network*: When all interactions where the user u_i directly mentions other users but the share of mentions in the Foreign Interactions set of all their SM Network users is very low, the posts authored by u_i are not very likely to have a great impact on his/her Network.
- *Low number of followers being rarely online*: Our impact metric rewards the users with a large network of active followers. If $\#SN(u_i)$ is low or $(ExposedUsers(u_i, E, \triangle t) \cap SN(u_i)^C$ is pretty high, the Impact metric is going to be low as well, as the number of users that can be impacted remains low.
- *Low social media activity location*: When the number of users in a location engaging with an Entity E is rather low, the impact metric can generate volatile results. In this circumstances it's advisable to extend the geographical coverage of the harvester (e.g.: increasing the radius).
- *High-activity SM location*: The design of the impact metric is not resolving overlappings in the *ExposedUsers* sets of the users behind the impact on an particular Entity E in a location. If the same SM user is part of different

ExposedUsers sets, the contributions of the impacting users rather add up, which intents to reflect the combined effect on the user being exposed.

5.2. *Design Decisions and performance evaluation*

To define our metrics and implement our system, several design decisions have been taken (e.g.: points we intentionally left unaddressed for further research, deliberate decisions against other approaches to solve punctual problems for the sake of simplicity, decisions where we opted for the most complete solution trading off simplicity for accuracy, etc). In this subsection, we go through each decision explaining the rationale behind it and pointing out alternatives for future research.

5.2.1. *Data Gathering*

- *Users geo-location*
  Our system relies on the geo-location capabilities of the Twitter Search API to periodically retrieve all the interactions of any kind created over SM channels in the specified area. A limitation is that some transactions created by users in the area are not geo-localized and can't therefore be retrieved by a geo-query. To overcome this problem, a potential solution would be implementing a *user-place stickiness factor*, which computes based on the user's history of interactions, the likelihood of a particular interaction to be located in the area under analysis. Implementing such an approach would improve the data gathering recall.

5.2.2. *Significance for the author*

- *Quantifying engagement*
  We defined the *Entity Engagement Index* as a share of *Entity Related Interactions* over the overall number of interactions to measure the relevance of the *Entity Related Interaction* within the set of all interactions authored by the user.
  Additionally, the system could separate interactions initiated by the user him/herself from forwarding behaviors (re-tweets, shares, etc, depending on the SM platform) to define a weighting schema based on the intensity —e.g.: a re-tweet would have lower weights than an interaction where the author is also the initiator of the conversational thread—.
  A point worth researching would also be understanding the effect of taking a share over the number of SM interactions in the same industry only (e.g.: Transportation)
- *Modeling the Differential Perception Factor*
  CARESOME considers all interactions in a category to have the same weight. Enhancing the purpose tagging with tonality —e.g.: based on sentiment analysis— and comparing it to the baseline tonality for the particular author can also be use to modulate the perception factor in future works.

5.2.3. *Communication Category Share Analysis*

- *Categories prevalence*
  CARESOME flags the set of harvested tweets in several waves (one per purpose category). If some terms from category A and other terms from category B have positive matching with the tweet content, a prevalence rule is triggered to decide which communication purpose category the tweet is assigned to (e.g.: the purpose with higher number of matching terms, or in case of a draw, the one defined as more dominant). The prevalence rules need to be defined and apply in the same way to all terms.
  In order to make these rules more effective, terms can be given a weight indicated how representative it is for the category (e.g.: for complaints, the bi-gram "no info" is less representative for "complaints" than "sucks"). Working with weights would help defining better which category the tweet should be flagged with.
- *Categories overlapping*
  Our implementation don't foresee that a given SM interaction can be flagged in 2 communication purpose categories (e.g.: "You lost my suitcase, now what? your service sucks" — *Information Request* and *Complaints*—). As the overlapping of purposes usual is, future research could implement a mechanism to address that.

- *Irony and Brand comparison*
  CARESOME doesn't implement any mechanism to handle irony. Taking into account all interactions from the author with the brand might help uncovering outliers (e.g.: many complaints and a punctual praise). Also situations where in a the same tweet 2 competiting entities are mentioned are not handled by CARESOME at present (e.g.: "After being in a @[Entity A] flight I cannot fly @[EntityB] anymore because the service sucks!").

5.2.4. Impact on the SM network

- *Defining Exposure*
  We worked with probability windows referred to the points of time where the user authored the interaction to engage with the Entity and the evidence that the user whose exposure is being checked has been active within this probability window. Our approach does not take into account users' activity patterns of any kind, likelihood of reading based on the time of the day where the interaction was created, etc. Exposure modeling is certainly a research line that can provide promising results with respect to certainly simplified way we have implemented it. In our system, we took the decision of defining a window, whereas more fuzzy-oriented implementations could have also been analyzed, like a decay-gradient function instead of the crisp simplification we implemented.
- *Defining Tie-Strength*
  Similarly, the interpretation for Tie-Strength we have implemented in our system might look simplistic. We opted for merely measuring the direct interactions —what we believe is the most defining factor— but keeping it bidirectional. Other factors might be thrown into the mix in further studies, like SM networks overlapping, Tie-Strength with common first degree connections, size of the SM network, etc.
- *Differential Influence*
  Putting Tie-Strength aside, the importance of the user within the social network of the follower has not been implemented for simplicity reasons. The impact caused by the interaction of certain user on another one depends to certain extent on how important the first one is within the SM network of the later. Ranking users within a SM network requires complex modeling which led us to postpone this aspect to further analysis.
  Another possible improvement would be introducing a reputation index for SM authors, which can be taken into account in the impact computation on the SM network. Another improvement could be achieved by modelling the quality of the interaction, approach which has delivered good results in the recommender system domain [66, 67, 68].

5.2.5. Optimizing for time-2-results

- *Computing Exposure*
  A good compromise between computing the Exposure at user level and applying a method to approximate it would be keeping memory creating a long-term index at user level. This long term index might also vary per user and per time of the day / day of the week, which adjusts much better to the interaction patterns in the SM world. Future work could significantly improve the approximation to the Exposure computation.
- *Tie-strength computing*
  Our suggestion to model the Tie-Strength is pretty simplistic and works reasonably well in scenarios where prompt decision taking is required. For those use cases where precision has priority over speed, the Tie-Strength can be redefined taking into consideration other factors like overlapping of SM Networks, interactions where both users are mentioned, etc.
  In general, the approximation of both Tie-Strength and Exposure as explained in 3.6.2 is per se a research area needing exploration.

6. Conclusions

In this paper we introduced CARESOME, a system that leverages geo-located SM insights to support both customer retention and acquisition activities. CARESOME turns the SM channels into a sensor that companies can use to understanding the impact of the unfiltered feedback given by their customers and prospect customers, but also to uncover competitors' weak spots and engineer acquisition strategies targeting them. Our system relies on a framework of metrics intended to quantify what we defined as intrinsic and extrinsic impact, where we modeled the contribution of all potential factors playing a role in the impact perception, such as author's engagement with the topic, the underlying communication purpose per interaction and how the authors of these interactions are connected to other SM users.

CARESOME is designed to produce actionable insights supporting the customer facing departments of any service company. Thus, in addition to the suggested approach to compute the impact metrics, a speed modus is available, which trades accuracy against time-to-results. To make the generated insights more actionable and enable a prompter decision making, CARESOME also implements a mapping of the results to categories so that the system users do not have to deal with large, hard to compare numbers, but with simple shaded impact categories over time.

To discuss the system performance we presented a real case scenario from the travel industry and engaged into a discussion about the design decisions, indicating potential limitations and pointing at further research lines to contribute to the evolution of CARESOME, especially in the impact modeling area.

Acknowledgements

This paper has been developed with the financing of projects TIN2010-17876, TIC5299, TIC-5991and TIN2013-40658-P.

References

[1] T. Hennig-Thurau, C. Wiertz, F. Feldhaus, 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) 1–32.

[2] A. D. Stein, M. F. Smith, R. A. Lancioni, The development and diffusion of customer relationship management (crm) intelligence in business-to-business environments, Industrial Marketing Management 42 (6) (2013) 855–861.

[3] M. Meadows, S. Dibb, Progress in customer relationship management adoption: a cross-sector study, Journal of Strategic Marketing 20 (4) (2012) 323–344.

[4] A. Sen, A. P. Sinha, It alignment strategies for customer relationship management, Decision Support Systems 51 (3) (2011) 609–619.

[5] Y. Chen, J. Xie, Online consumer review: Word-of-mouth as a new element of marketing communication mix, Management Science 54 (3) (2008) 477–491.

[6] D. Kempe, J. Kleinberg, É. Tardos, Maximizing the spread of influence through a social network, in: Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining, ACM, 2003, pp. 137–146.

[7] D. Yates, S. Paquette, Emergency knowledge management and social media technologies: A case study of the 2010 haitian earthquake, International Journal of Information Management 31 (1) (2011) 6–13.

[8] T. Sakaki, M. Okazaki, Y. Matsuo, Earthquake shakes twitter users: Real-time event detection by social sensors, in: Proceedings of the 19th International Conference on World Wide Web, WWW '10, ACM, New York, NY, USA, 2010, pp. 851–860.

[9] H. Gao, G. Barbier, R. Goolsby, Harnessing the crowdsourcing power of social media for disaster relief, IEEE Intelligent Systems 26 (3) (2011) 10–14.

[10] C. D. Corley, D. J. Cook, A. R. Mikler, K. P. Singh, Text and structural data mining of influenza mentions in web and social media, International journal of environmental research and public health 7 (2) (2010) 596–615.

[11] M. Thomson, F. Doblas-Reyes, S. Mason, R. Hagedorn, S. Connor, T. Phindela, A. Morse, T. Palmer, Malaria early warnings based on seasonal climate forecasts from multi-model ensembles, Nature 439 (7076) (2006) 576–579.

[12] R. Basher, Global early warning systems for natural hazards: systematic and people-centred, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 364 (1845) (2006) 2167–2182.

[13] X. Zhang, H. Fuehres, P. A. Gloor, Predicting stock market indicators through twitter "i hope it is not as bad as i fear", Procedia-Social and Behavioral Sciences 26 (2011) 55–62.

[14] B. Candelon, E.-I. Dumitrescu, C. Hurlin, How to evaluate an early-warning system: Toward a unified statistical framework for assessing financial crises forecasting methods, IMF Economic Review 60 (1) (2012) 75–113.

[15] J. Smailović, M. Grčar, N. Lavrač, M. Žnidaršič, Stream-based active learning for sentiment analysis in the financial domain, Information Sciences (2014) (in press).

[16] M. Bussière, M. Fratzscher, Towards a new early warning system of financial crises, journal of International Money and Finance 25 (6) (2006) 953–973.

[17] C. Jiang, K. Liang, H. Chen, Y. Ding, Analyzing market performance via social media: a case study of a banking industry crisis, Science China Information Sciences 57 (5) (2014) 1–18.

[18] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. Van Nes, M. Rietkerk, G. Sugihara, Early-warning signals for critical transitions, Nature 461 (7260) (2009) 53–59.

[19] T. Hennig-Thurau, E. C. Malthouse, C. Friege, S. Gensler, L. Lobschat, A. Rangaswamy, B. Skiera, The impact of new media on customer relationships, Journal of Service Research 13 (3) (2010) 311–330.

[20] C. Chen, C.-Y. Zhang, Data-intensive applications, challenges, techniques and technologies: A survey on big data, Information Sciences (2014) (in press).

[21] A. M. Kaplan, M. Haenlein, Users of the world, unite! the challenges and opportunities of social media, Business horizons 53 (1) (2010) 59–68.
[22] W. G. Mangold, D. J. Faulds, Social media: The new hybrid element of the promotion mix, Business horizons 52 (4) (2009) 357–365.
[23] B. J. Jansen, M. Zhang, K. Sobel, A. Chowdury, Twitter power: Tweets as electronic word of mouth, Journal of the American society for information science and technology 60 (11) (2009) 2169–2188.
[24] T. Li, G. Berens, M. de Maertelaere, Corporate twitter channels: The impact of engagement and informedness on corporate reputation, International Journal of Electronic Commerce 18 (2) (2013) 97–126.
[25] A. Java, X. Song, T. Finin, B. Tseng, 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, pp. 56–65.
[26] Y. Kim, M. A. Ahmad, Trust, distrust and lack of confidence of users in online social media-sharing communities, Knowledge-Based Systems 37 (2013) 438–450.
[27] J. Marés, V. Torra, On the protection of social networks users information, Knowledge-Based Systems 49 (2013) 134–144.
[28] M. J. Culnan, P. J. McHugh, J. I. Zubillaga, How large us companies can use twitter and other social media to gain business value., MIS Quarterly Executive 9 (4) (2010) 112–115.
[29] J. M. Chan, R. Yazdanifard, How social media marketing can influence the profitability of an online company from a consumer point of view, Journal of Research in Marketing 2 (2) (2014) 157–160.
[30] A. Rapp, L. Beitelspacher, D. Grewal, D. Hughes, Understanding social media effects across seller, retailer, and consumer interactions, Journal of the Academy of Marketing Science 41 (5) (2013) 547–566.
[31] J. A. Chevalier, D. Mayzlin, The effect of word of mouth on sales: Online book reviews, Journal of marketing research 43 (3) (2006) 345–354.
[32] J. Villanueva, S. Yoo, D. M. Hanssens, The impact of marketing-induced versus word-of-mouth customer acquisition on customer equity growth, Journal of marketing Research 45 (1) (2008) 48–59.
[33] R. N. Bolton, A dynamic model of the duration of the customer's relationship with a continuous service provider: the role of satisfaction, Marketing science 17 (1) (1998) 45–65.
[34] R. Rishika, A. Kumar, R. Janakiraman, R. Bezawada, The effect of customers' social media participation on customer visit frequency and profitability: an empirical investigation, Information systems research 24 (1) (2013) 108–127.
[35] N. Naveed, T. Gottron, J. Kunegis, A. C. Alhadi, Bad news travel fast: A content-based analysis of interestingness on twitter, in: Proceedings of the 3rd International Web Science Conference, ACM, 2011, p. 8.
[36] J. Park, M. Cha, H. Kim, J. Jeong, Managing bad news in social media: A case study on domino's pizza crisis., in: ICWSM, 2012.
[37] T. Sakaki, M. Okazaki, Y. Matsuo, Earthquake shakes twitter users: real-time event detection by social sensors, in: Proceedings of the 19th international conference on World wide web, ACM, 2010, pp. 851–860.
[38] A. Culotta, Towards detecting influenza epidemics by analyzing twitter messages, in: Proceedings of the first workshop on social media analytics, ACM, 2010, pp. 115–122.
[39] S. Middleton, L. Middleton, S. Modafferi, Real-time crisis mapping of natural disasters using social media, IEEE Intelligent Systems 99 (PrePrints) (2014) (in press).
[40] J. Yin, A. Lampert, M. Cameron, B. Robinson, R. Power, Using social media to enhance emergency situation awareness, IEEE Intelligent Systems 27 (6) (2012) 52–59.
[41] R. Colbaugh, K. Glass, Early warning analysis for social diffusion events, Security Informatics 1 (1) (2012) 1–26.
[42] N. R. Adam, B. Shafiq, R. Staffin, Spatial computing and social media in the context of disaster management, IEEE Intelligent Systems 27 (6) (2012) 90–96.
[43] C. Dellarocas, The digitization of word of mouth: Promise and challenges of online feedback mechanisms, Management science 49 (10) (2003) 1407–1424.
[44] Y. Liu, Word of mouth for movies: Its dynamics and impact on box office revenue, Journal of marketing 70 (3) (2006) 74–89.
[45] N. I. Bruce, N. Z. Foutz, C. Kolsarici, Dynamic effectiveness of advertising and word of mouth in sequential distribution of new products, Journal of Marketing Research 49 (4) (2012) 469–486.
[46] E. Bothos, D. Apostolou, G. Mentzas, Using social media to predict future events with agent-based markets, IEEE Intelligent Systems 25 (6) (2010) 50–58.
[47] H. Kwak, C. Lee, H. Park, S. Moon, What is twitter, a social network or a news media?, in: Proceedings of the 19th International Conference on World Wide Web, WWW '10, ACM, New York, NY, USA, 2010, pp. 591–600.
[48] L. Page, S. Brin, R. Motwani, T. Winograd, The pagerank citation ranking: Bringing order to the web, in: Proceedings of the 7th International World Wide Web Conference, Brisbane, Australia, 1998, pp. 161–172.
[49] S. Ye, S. F. Wu, Measuring message propagation and social influence on twitter.com, in: Proceedings of the Second International Conference on Social Informatics, SocInfo'10, Springer-Verlag, Berlin, Heidelberg, 2010, pp. 216–231.
[50] M. Cha, H. Haddadi, F. Benevenuto, K. P. Gummadi, Measuring user influence in twitter: The million follower fallacy, in: ICWSM10: Proceedings of international AAAI Conference on Weblogs and Social, 2010.
[51] D. M. Romero, B. Meeder, J. Kleinberg, 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, WWW '11, ACM, New York, NY, USA, 2011, pp. 695–704.
[52] C. C. Yang, X. Tang, Estimating user influence in the medhelp social network, IEEE Intelligent Systems 27 (5) (2012) 44–50.
[53] P. V. Marsden, K. E. Campbell, Measuring tie strength, Social forces 63 (2) (1984) 482–501.
[54] M. Granovetter, The strength of weak ties, American journal of sociology 78 (6) (1973) l.
[55] E. Gilbert, K. Karahalios, Predicting tie strength with social media, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM, 2009, pp. 211–220.
[56] C. Haythornthwaite, Tie strength and the impact of new media, in: Proceedings of the annual Hawaii International Conference on System Sciences, 2001, p. 11.
[57] P. A. Grabowicz, J. J. Ramasco, E. Moro, J. M. Pujol, V. M. Eguiluz, Social features of online networks: The strength of intermediary ties in

online social media, PloS (Public Library of Science) one 7 (1) (2012) 1–27.
[58] R. K. Pan, J. Saramäki, The strength of strong ties in scientific collaboration networks, EPL (Euro Physics Letters) 97 (1) (2012) 18007.
[59] H. Shin, J. Lee, Impact and degree of user sociability in social media, Information Sciences 196 (2012) 28–46.
[60] R. Likert, A technique for the measurement of attitudes, Archives of psychology 22 (140) (1932) 1–55.
[61] Q. X. Yang, S. S. Yuan, L. Zhao, L. Chun, S. Peng, Faster algorithm of string comparison, Pattern Analysis & Applications 6 (2) (2003) 122–133.
[62] M. Krieger, D. Ahn, Tweetmotif: Exploratory search and topic summarization for twitter., 2010.
[63] T. Kiss, J. Strunk, Unsupervised multilingual sentence boundary detection, Computational Linguistics 32 (4) (2006) 485–525.
[64] X. Wang, A. McCallum, X. Wei, Topical n-grams: Phrase and topic discovery, with an application to information retrieval, in: Data Mining, 2007. ICDM 2007. Seventh IEEE International Conference on, IEEE, 2007, pp. 697–702.
[65] B. C. Gencosman, H. C. Ozmutlu, S. Ozmutlu, Character n-gram application for automatic new topic identification, Information Processing & Management 50 (6) (2014) 821–856.
[66] J. Serrano-Guerrero, E. Herrera-Viedma, J. A. Olivas, A. Cerezo, F. P. Romero, A google wave-based fuzzy recommender system to disseminate information in university digital libraries 2.0, Information Sciences 181 (9) (2011) 1503–1516.
[67] C. Porcel, T.-L. A., M. Martnez, E. Herrera-Viedma, A hybrid recommender system for the selective dissemination of research resources in a technology transfer office, Information Science 184 (1) (2012) 1–19.
[68] A. Tejeda-Lorente, C. Porcel, E. Peis, R. Sanz, E. Herrera-Viedma, A quality based recommender system to disseminate information in a university digital library, Information Science 261 (2014) 52–69.

2.2 Leveraging Localized Social Media Insights for Industry Early Warning Systems

- Bernabé-Moreno, J., A. Tejeda-Lorente, C. Porcel and E. Herrera-Viedma. "Leveraging Localized Social Media Insights for Industry Early Warning Systems". International Journal of Information Technology and Decision Making (2015).
  - Status: **Accepted for Publication**.
  - Impact Factor (JCR 2013): 1.890
  - Subject Category: Operation Research and Management Science. Ranking 12 / 79 (**Q1**).
  - Subject Category: Computer Science, Information Systems. Ranking 24 /135 (**Q1**).

International Journal of Information Technology & Decision Making

LEVERAGING LOCALIZED SOCIAL MEDIA INSIGTHS FOR INDUSTRY EARLY WARNING SYSTEMS

JUAN BERNABÉ-MORENO

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

ÁLVARO TEJEDA-LORENTE

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

CARLOS PORCEL-GALLEGO

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

ENRIQUE HERRERA-VIEDMA

Department of Computer Science and A.I., University of Granada
Granada E-18071, Spain Country[*]

Social Media has become the easiest, cheapest and fastest channel for companies to identify the events that affect their customers. The geo-location capabilities of the Social Media interactions, enables Early Warning systems to alert not only when the quality of service decays, but also where and how many customers are impacted. In this paper we present a system based on geo-localized Social Media and the corresponding 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.

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

1. Introduction

In the field of disaster prevention, Early Warning Systems (EWS) have been extensively developed in different scenarios, such as pandemic expansion [1,2], earthquakes [3,4], flood and other natural hazards [5], etc. Alerting about events *as they happen* has also proven very useful in the industry domain, where EWS have been implemented to gain a competitive advantage over all competitors, suppliers and other market

[*]Department of Computer Science and Artificial Intelligence, University of Granada, Granada E-18071, Spain. viedma@decsai.ugr.es.

players. In the financial domain, EWS, fed by a large variety of economic indicators, have been employed to assess at a macro level the vulnerability of the markets [6], to detect financial crisis at early stages [7,8] or at a much finer granular level, to identify critical transactions [9], etc.

The number of companies providing services instead of one-off products is drastically increasing, especially with the advent of cloud technologies and the emerging of the *as-a-service* business models [10]. For a company, monitoring the quality of service offered across all locations where its customers are active and acting as soon as possible when a quality decay is registered is becoming a must-do to survive in competitive market. The first step is to understand the impact of the potential service disruption on the set of affected customers; once the impact has been quantified, the way of handling the incident is defined -which might even be, not doing anything-; as a third step and if there's a need to take action, a reaction plan is put in place and executed to prevent customers churn. Therefore, EWS have become an integral part of the service providers operations to implement customer retention strategies [11].

The ever increasing adoption of Social Media (SM) unlocks new possibilities for implementing EWS based on customers' feedback. The nature of the customers-brand interaction over the SM channel presents 2 characteristics, that make SM highly suitable for implementing EWS on top: it's *near-real time* -just a bit of delay- and it's *unfiltered* -no censorship on it-. Therefore, more and more service providers take advantage of the SM channel as an extension of their customer care activities, but also as an additional channel to run customers' acquisition campaigns. SM platforms are usually open, which likewise allows for setting up monitors for not only own customers but also for competitors' customers: EWS could alert on a service quality decay of a competitor, so that the campaign machinery starts targeting these foreign customers to make them churn from the competitors in favour of the own proposition.

SM experimented a new (maybe minor) revolution when the access to the Internet escaped the confines of the desktop and became mobile. When network technologies made the mobile data transfer possible, SM platforms capitalized on these new capabilities and became more pervasive. New functionality was developed so that with the user's consent, to each and every SM interaction a place-stamp -in form of a pair of geographical coordinates, the name of a place, etc- could be added. The combination *time-stamp* and *place-stamp* introduces new analytical possibilities that were not feasible before. Those companies, that have been analyzing the SM channels to address different use cases at a broader geographical level –i.e.: 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.

The location based insights generated out of the SM channels are invaluable inputs to source into EWS of any kind, which then can also produce alerts for locations in particular (as shown in Fig. 1)

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.

This paper has been structured starting with the background for our research in the Section 2). The impact metrics are introduced in Section 3. In the Section 4 we then present the system implemented to create the metrics and in the Section 5 we discuss the results obtained for a real world scenario. Finally, we closure off the paper throwing our conclusions and pointing to further research lines.

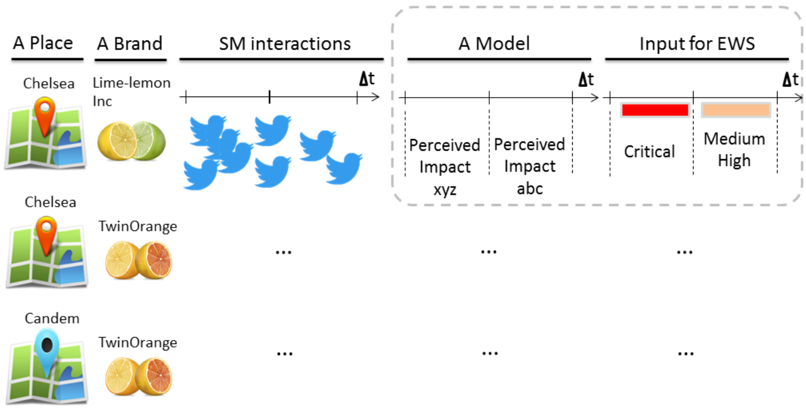

Fig. 1. Geo-localized Social Media as Input for Early Warning Systems

2. Background

The extensive use of SM by companies as a channel to obtain near-real time uncensored feedback from their customers turn the SM adoption aspect into a very active and literature rich research domain. We point out the seminal paper authored by Kaplan et al. [12], where for the first time, the SM integration with the traditional media to reach customers was declared as a must and SM labelled as the channel to engage with customers in a time-close and high-efficient manner. It was in [13] where the authors first time considered SM as integral part of the promotion mix highlighting the less controlled and hence more insights revealing nature of customers interactions.

The majority of the research papers focused in one way or another on Twitter as SM platform, because they grant the access to much more information in a very structured and effective way offering different APIs -unlike other competing platforms such as Facebook or LinkedIn-. 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. In [14] for example, a thorough analysis on how all the interactions created over Twitter about a brand impacts its corporate image. Hennig et. al, in [15] explained how microblogging was shaking traditional business models by making more crucial the attention to product quality, as with platforms like Twitter, the time elapsed between product launch and first feedback available shrank dramatically. In [16] the authors explore the relation between intentions and connectivity in the same customer based.

The speed of diffusion in SM has been also subject of countless pieces of research. As explained in [17] and [18], 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. In [19], 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 in [20] a method to detect epidemic expansion on early stages. In [21], a method is proposed to detect local events in a real-time manner based on the Twitter stream. Rosi et al. in [22] provides an overview of the applications of social sensing in pervasive environments. In [23] 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. Kwak [24] defined 3 metrics aimed at quantifying the *social influence*: the so called *propagation influence*, based on the Google Search PageRank algorithm [25], *followers influence* –more followers implies more influence–, and *re-tweet influence* –more re-tweets means more influence–. Ye and Wu [26] 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–. Cha [27] 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. [28] develop a mechanism to quantify how the exposure to other users is making them adopt a new behavior. In the framework we are presenting in this article, we evolved the concept of exposure by defining different contribution levels to the total impact of the topic.

One of our metrics is based on the polarity of the SM interaction content. Sentiment Analysis applied to SM has been subject of prolific research. The ground work derives from all the previous studies of terms polarity in the Natural Language Processing domain [29,30]. Pang et al. in [31] set the basis of opinion mining based on the analysis of sentiments. Kouloumpis et al. in [32] and Agarwal et al. in [33] provided an extensive research on sentiment analysis applied to microblogging messages. In this work we rely on SentiWordNet 3.0, implemented by Baccionella et al. in [34] on top of the WordNet lexical English database [35]

3. Framework definition

The ultimate aim of our framework is providing a way 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 modelled by reach –or number of exposed users to this interaction–, a *severity level* to express how critical the situation is 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 and to feed EWS. 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 are built upon:

Definition 3.1. The set U represents the set of Social Media Users from which we have evidence they have been in the location L ($InLocation(u_i, L, \triangle t)$) we are monitoring during the time period under analysis $\triangle t$

$$U \equiv \{u_i\}\ (i = 1, ..., n),\ InLocation(u_i, L, \triangle t) \tag{1}$$

Definition 3.2. The Social Network for a given user u_i is defined as:

$$SN\ (u_i) \equiv \{u_j\}(j = 1, ..., n),\ \forall\ u_j \in SN(u_i),\ Follows(u_i, u_j) \tag{2}$$

$Follows(u_i, u_j)$ is a relation 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 $Follows(u_i, u_j) \not\Leftrightarrow Follows(u_j, u_i)$

The fact that a user u_j is part of the SN of another user u_i does not necessarily mean that u_j has to be located in the same location L of user u_i: $u_j \in SN(u_i), u_i \in U \not\Rightarrow u_j \in U$, as $u_j \in SN(u_i), u_i \in U \not\Rightarrow InLocation(u_j, L, \triangle t)$

Definition 3.3. The set $SN(U)$ represents the set of all the users being followed by the users in U:

$$SN(U) \equiv \cup_{i=1}^{|U|} SN(u_i) \tag{3}$$

Definition 3.4. We define all user interactions ($UserInteractions$) for a given user u_i over a time interval $\triangle t$, as:

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

Definition 3.5. A Social Media Interaction it represents the atomic piece of content authored by the user u_i during the time $\triangle t$ in a Social Media Platform (e.g.: a tweet, a re-tweet).

The function $Author(u_i, it_i, \triangle t)$ returns $True$ if u_i created the interaction it_i in the time period $\triangle t$, and $False$ otherwise.

The time interval $\triangle t$ 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 3.6. 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 3.7. We define all user interactions ($Interactions$) for a given user u_i over a time interval $\triangle t$, as:

$$Interactions\ (u_i, \triangle t) \equiv \{it_i\}(i = 1, ..., n),\ \forall it_i \in Interactions\ (u_i, \triangle t), \\ Author(u_i, it_i, \triangle t) \tag{5}$$

Definition 3.8. We define the set of $Interactions$ for a given user u_i with the entity E over a time interval $\triangle t$ as:

$$Interactions\ (u_i, E, \triangle t) \equiv \{it\},\ \forall it_i \in Interactions\ (u_i, E, \triangle t),\ Author(u_i, it_i, \triangle t) \\ \wedge related(it_i, E) \tag{6}$$

Where $related(it_i, E)$ is a NLP membership function retrieving $True$ if the interation 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:

$$Engaged(u_i, E, \triangle t) \equiv True, \exists it_i,\ it_i \in Interactions(u_i, E, \triangle t), u_i \in U \cup SN(U) \quad (7)$$

Where u_i is the user, E is the representation of the Entity, $\triangle t$ is the time span specified consisting of two components (*t_startdate* and *t_enddate*), it_i represents a social media interaction and $Interactions\ (u_i, E, \triangle t)$ represents the interactions of the user u_i related to the Entity E in the time interval $\triangle t$, 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, the so called *Entity Engagement Index (EEI)*:

$$EEI(u_i, E, \triangle t) = \frac{|Interactions\ (u_i, E, \triangle t)|}{|\cup_{k=1}^{|E|} Interactions\ (u_i, E_k, \triangle t)|} \quad (8)$$

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

The Entity Engagement Index can also be expressed as a share of the interactions related to one entity over all interactions:

$$EEI(u_i, E, \triangle t) = \frac{|Interactions\ (u_i, E, \triangle t)|}{|Interactions(u_i, \triangle t)|} \quad (9)$$

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 3.9. 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 a particular communicative intention (such as praise, criticism, information inquiry, complaints, etc).

Even if the boundaries might not be crisp, we can assign each interaction to a *leading Purpose Category* within the set of purpose categories considered PC:

$$\forall it_i \in Interactions\ (u_i, E, \triangle t), \exists p_k, Purpose(it_i) = p_k, p_k \in PC \quad (10)$$

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

$Interactions(u_i, P, \triangle t)$ represents the set of all interactions authored by a user u_i over the period of time $\triangle t$ whose leading Purpose Category is P.

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:

$$Interactions\ (u_i, E, P, \triangle t) = Interactions\ (u_i, E, \triangle t) \cap Interactions\ (u_i, P, \triangle t) \tag{11}$$

$$DPF(u_i, E, P, \triangle t) = \frac{|Interactions\ (u_i, E, P, \triangle t)|}{|Interactions\ (u_i, P, \triangle t)|} \tag{12}$$

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.

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

$$Exposed\ (u_i, u_j, E, \triangle t) = \begin{cases} True, & u_j \in SN(u_i), \exists it_k, it_k \in Interactions(u_i, E, \triangle t), \\ & P(read(u_j, it_k, \triangle t)) \geq Threshold \\ False, & otherwise \end{cases} \tag{13}$$

where $P(read(u_j, it_k, \triangle t))$ is the probability that the user u_j reads the content posted in the interaction it_j in the designated time $\triangle t$. 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:

$$ExposedUsers(u_i, E, \triangle t)\ \equiv \{u\},\ \forall u_j,\ Exposed(u_i, u_j, E, \triangle t)\ =\ True,\ u_i \in U \tag{14}$$

Apart from modelling how many users are exposed to the Interactions with a given communication purpose and modulating it by the aforementioned *DPF*, we need a way of understanding the level of severity of the incident reported by the users in the location.

As a severity indicator, we suggest a set of metrics based on the positive and negative polarities of each and every SM interaction message in a place. The first component is what we called the *Severity Value*, defined as follows:

$$SeverityValue(u_i, E, P, \triangle t) = \sum_{i=1}^{|Int(u_i,E,P,\triangle t)|} \frac{w\_pos * Pos(it_i) + w\_neg * Neg(it_i)}{|Interactions\ (u_i, E, P, \triangle t)|} \tag{15}$$

After that we establish a set of *Severity Levels* applying a mapping function to obtain meaningful and workable levels of severity:

$$SeverityLevel(u_i, E, P, \triangle t) = Map2Level(SeverityValue(u_i, E, P, \triangle t)) \tag{16}$$

Where $Pos(it_i) \in \mathbb{R}$ and $Neg(it_i) \in \mathbb{R}$ are the positive and negative polarity values of the SM interaction content it_i, w_pos and w_neg represent the weights given to the positive and negative polarity values. For example if the Purpose Category is *"Complaint"*, w_pos value might be chosen to be 0 and w_neg to be 1, as just negative polarity is meaningful for the severity modelling in this case. If the Category was *"Praise"*, the values of w_pos and w_neg would be the other way around, or for a Category like *"Information Request"*, we might choose both to be 1, as both are equally significant. $Map2Level(x)$ is a function mapping a polarity value to a Severity Level $Sin(1, .., n)$ being n the number of levels predefined.

In order to make the Severity Level usable in the computation of the Impact metric, we map it to a value which is going to serve as a contribution factor for multiplying by the other components in the final Impact equation, as we are going to present later in this section:

$$SF(u_i, E, P, \triangle t) = Map2Factor(SeverityLevel(u_i, E, P, \triangle t)) \tag{17}$$

Where $Map2Factor(SeverityLevel(u_i, E, P, \triangle t))$ assigns a value between 0 and 1 to the different Severity Levels (e.g.: if we considered 4 levels, the level 1 would map to 1, the level 2 to 0.75, the 3rd one to 0.5 and the last one to 0.25).

Even if for simplicity the mapping of polarity values to SF could be done without having to do the intermediate step of mapping to the Severity Levels, it's good to have them as a control step. In Table 3.4 we provide one example to explain how the Severity related metrics work: the negative and positive polarities are obtained for 4 SM interactions; the weightings are then applied to both positive and negative polarity values and the function *Map2Level* assigns after that a level to each interaction. In a last step, the function *Map2Factor* assigns the corresponding values to the previously obtained levels, which is what goes in the Impact equation, as we are going to see below.

			w_pos=0,5	w_neg=0,5	Difference	Map2Level	Map2Factor
Interaction	Pos	Neg	w_pos*Pos	w_neg*Neg	SevValue	SevLevel	SF
it1	0,25	1,25	0,13	0,63	-0,5	2	0,75
it2	1,25	1,50	0,63	0,75	-0,125	3	0,5
it3	0,75	2,50	0,38	1,25	-0,875	1	1
it4	0,00	0,75	0,00	0,38	-0,375	3	0,5

Table 1. Example of Severity Computation

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

$$\begin{aligned} Impact(u_i, E, P, \triangle t) = \Im(&EEI(u_i, E, \triangle t), \\ &DPF(u_i, E, P, \triangle t), \\ &\#ExposedUsers(u_i, E, \triangle t), \\ &SF(u_i, E, P, \triangle t)) \end{aligned} \tag{18}$$

The $\Im$ 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:

$$DPF(u_i, E, P, \triangle t) \approx 1 \tag{19}$$

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:

$$\#ExposedUsers(u_i, \triangle t) \;\approx \#SN(u_i) * K, \;\; K \in [0, 1] \tag{20}$$

3.5. *Consolidated Impact Metric*

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

$$Impact(U, E, P, \triangle t) = \sum_{i=1}^{|U|} Impact(u_i, E, P, \triangle t) \tag{21}$$

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. The number of levels depends on the particular use of this information by the EWS later on: if the EWS provides different alerting levels driving to different action plans, these levels should be considered here.

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 2 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.

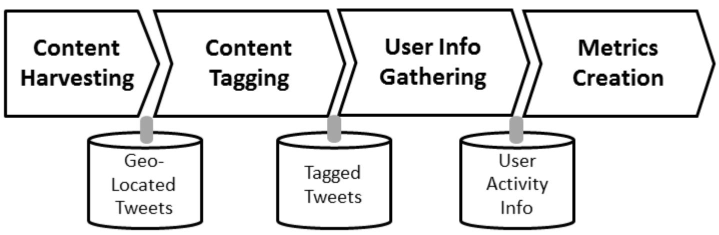

Fig. 2. System structure

4.1. *Content Harvesting*

Relying on the Twitter Search API publicly available Twitter Search API [a] 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. This module performs poll-requests from the Twitter Search API to store the tweets into a local data base for further processing. The tweets are selectively picked for a given area which is configured in the harvester, namely the one we want to perform the topic impact analysis over time. Additionally, the Twitter API supports the filtering by language (e.g.: only tweets in English), but even it would make the later NLP much easier, it might disregard the interactions of all users related with the topic in the target area for being in a foreign language. We opted for a work-around that doesn't filter out the tweets by language upfront, yet doesn't introduce the need for applying NLP techniques in all identified languages, as we are going to explain in the next section.

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 function [36] to accept spelling mistakes (e.g.: *lutfhansa* 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 [37,38] over a long history of content related to a particular category,

[a]Available at *https://dev.twitter.com/docs/api/1/get/search?*

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 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 19 and 20 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 $ExposedUsers(u_i, \triangle t)$ 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 $it_j \in Interactions(u_j, \triangle t)$ 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, $\exists it_i \in Interactions(u_i, E, \triangle t), [t(it_j - 60min, t(it_j + 60min] \cap [t(it_j - 60min, t(it_j + 60min] \neq \oslash$. 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 $\triangle t$ 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 (18) is computed. It requires the previous calculation of the single components: the Entity Engagement Index (equation 9), the DPF (equation 12) and the size of the Exposed Users (equation 14). For all the interactions created by the SM user in the location under analysis related to the Entity and with the right leading Communication Purpose, the Severity Values are computed (equation 15) and the the Severity Levels (equation 15) that are then mapped to produce the Severity Factors (equation 17) required in the Impact equation (18). Once the individual Impact has been computed, the overall Impact is obtained applying the aggregation (eq. 21).

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

We have chosen a real world scenario to demonstrate how our system and the impact metrics can leverage SM insights to feed EWS. Our experiment is based on two harvesters located in the main airports in the city of London: Heathrow and Gatwick (having the center placed right in the middle of the respective airports and a radius of 5 km). We let the harvesters run for exactly 2 months, from the 23^{rd} of November 2013 to 23^{rd} of January 2014, gathering a total of 852319 SM interactions (more precisely tweets). During this period of time there were severe weather conditions, spreading the chaos all over the country with strong winds and flooding episodes, so we could detect quality of service decay in several industries, especially transportation.

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 pushes people to express their discontent over the SM channels.

We considered Virgin Trains [b], First Capital Connect [c], National Rail [d], the companies offering exclusive express services Gatwick Express [e] and Heathrow Express [f] and the local operator Southern[g]

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 3 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*.

To model the severity in accordance with the selected Communication Purpose, we just considered the negative polarity (weighting $w\_pos = 0$ and $w\_neg = 1$ in the equation 15). The $Map2Level$ function (equation 16) and $Map2Factor$ (equation 17) are respectively given in Table 5.

Similarly, the Entities have been modelled including all relevant account information, hashed tags and even non-official accounts, like *@Southern_Trains* created

[b] `http://www.virgintrains.co.uk/`
[c] `http://www.firstcapitalconnect.co.uk/`
[d] `http://www.nationalrail.co.uk/`
[e] `http://www.gatwickexpress.com/`
[f] `https://www.heathrowexpress.com/`
[g] `http://www.southernrailway.com/`

Polarity	Level	Factor
[1,)	1	1
[0.75,1)	2	0.95
[0.5,0.75)	3	0.90
[0.25,0.5)	4	0.85
[0,0.25)	5	0.70

Table 2. Severity Levels and Factors

Top 10	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%
Top 20	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. 3. Penetration of the top 20 words categorizing a complaint across carriers

as a parody of the official *@SouthernRailUK*. In Figure 4 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 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
Compaints	5911	4	545	4779	352	231	0
Heathrow Airport							
Total	2817	26	17	2430	9	0	335
Compaints	668	8	10	572	2	0	75

Fig. 4. Number of Interactions per Twitter Harvester in Total and identified as a Complaint

In Figure 5 we can see the average Severity Level computed over time with the before specified values and functions for all companies in the different airports. Obviously for companies with a few SM interactions, the definition of Severity Levels is quite volatile (like Virgin in both airports or FCC in Heathrow). When the number of SM interactions is significant, we observe quite high severity levels which fits very well with the kind of engagement between users and transport companies (quite common strong negative opinion all along the period of time we analyzed).

Figure 6 shows the Impact metric computed for all 6 entities over 2 months in both airports. The highest value originated on Dec 24 for the Entity *National Railway*, reflects the train service disruption when the storm *Emily* was striking the country [h]. The second highest registered also for National Railway on Jan 17 was due again to weather causing flooding [i]. 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 [j]. Our metric can also quantify the impact of small decays in the quality of service. Figure 7 a) 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 7 b) –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 EWS 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.

Our metrics are designed to work in finer levels of granularity, whenever the number of SM interactions remains at a significant level. For the National Railway we have created an hourly heat-map (see Fig. 8), showing the value of hour impact assigned to deciles that we shaded according to the intensity by hour. In *a* we show all the values, but in *b* we just shaded the top decile pointing out the cases where the EWS shall generate an alert to trigger a triaging plan.

6. Conclusions

In this work we present a set of new metrics to measure how a community of customers located in a place is impacted when the quality of the service provided by a company decays and to quantify to which extent the company's image is affected. These metrics are built upon the SM interactions related to the company that are created in that particular location and address several aspects, such as the underlying communication purpose of the interaction, how the authors of these interactions are connected to other SM users and the level of severity, which we computed based on the content polarity.

[h] Southern services suspended due to strong storm on December 24th `http://www.bbc.com/news/uk-25785804`

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

[j] `http://www.londontoolkit.com/blog/daytrips/london-tours-on-christmas-day-new-years-2012/`

Our metrics are designed to produce insights to be fed into Early Warning Systems for decision making. Therefore, we provide an abstraction layer on top of the metrics mapping their values to levels that can be rapidly actioned by EWS. We also address the cases where the time to results is critical by providing approximations to the single components of our metric and removing therefore the time consuming steps but trading some precision off.

In our approach and for simplicity reasons we didn't consider the tie-strength between the author of a SM interaction and his/her network. Also for simplicity reason we didn't include in the severity definition the polarity baseline of the author. These 2 research lines might bring more precision to our metrics in future works. A certainly interesting area with a lot of potential especially in the field of Early Warning Systems would be exploring the usage of predictive capabilities on top of our framework.

Acknowledgements

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

References

1. Bouma MJ and Kaay HJ (1996). *The EI Nio 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", Vol. 1(1). 86-96.
2. Thomson M, Doblas-Reyes F, Mason S, Hagedorn R, Connor S, Phindela T, Morse A and Palmer T (2006). *Malaria early warnings based on seasonal climate forecasts from multi-model ensembles.* "Nature", Vol. 439(7076). 576-579
3. Allen RM and Kanamori H (2003). *The potential for earthquake early warning in southern California.* "Science", Vol. 300(5620).
4. Wu Y and Teng T (2002). *A virtual subnetwork approach to earthquake early warning.* "Bulletin of the Seismological Society of America", Vol. 92(5).
5. Basher R (2006). *Global early warning systems for natural hazards: systematic and people-centred.* "Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences", Vol. 364(1845).
6. Goldstein M, Kaminsky GL and Reinhart CM (2000). "Assessing financial vulnerability: an early warning system for emerging markets". Peterson Institute.
7. Bussire M and Fratzscher M (2006). *Towards a new early warning system of financial crises.* "journal of International Money and Finance", Vol. 25(6).
8. Edison HJ (2000). *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".
9. Scheffer M, Bascompte J, Brock WA, Brovkin V, Carpenter SR, Dakos V, Held H, Nes EHV, Rietkerk M and Sugihara G (2009). *Early-warning signals for critical transitions.* "Nature", Vol. 461(7260).
10. Chang WY, Abu-Amara H and Sanford JF (2010). "Transforming enterprise cloud services". Springer.
11. Reichheld, F. F., Sasser Jr, W. E. (1989). *Zero defections: quality comes to services.* "Harvard business review", 68(5), 105-111.

12. Kaplan AM and Haenlein M (2010). *Users of the world, unite! The challenges and opportunities of Social Media.* "Business horizons", Vol. 53(1).
13. Mangold WG and Faulds DJ (2009). *Social media: The new hybrid element of the promotion mix.* "Business horizons", Vol. 52(4).
14. Jansen BJ, Zhang M, Sobel K and Chowdury A (2009). *Twitter power: Tweets as electronic word of mouth.* "Journal of the American society for information science and technology", Vol. 60(11).
15. Hennig-Thurau T, Wiertz C and Feldhaus F (2012). *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".
16. Java A, Song X, Finin T and Tseng B (2007). *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.
17. Naveed N, Gottron T, Kunegis J and Alhadi AC (2011). *Bad news travel fast: A content-based analysis of interestingness on twitter.* In "Proceedings of the 3rd International Web Science Conference". ACM.
18. Park J, Cha M, Kim H and Jeong J (2012). *Managing Bad News in Social Media: A Case Study on Domino's Pizza Crisis.* In "ICWSM".
19. Sakaki T, Okazaki M and Matsuo Y (2010). *Earthquake shakes Twitter users: real-time event detection by social sensors.* In "Proceedings of the 19th international conference on World wide web". ACM.
20. Culotta A (2010). *Towards detecting influenza epidemics by analyzing Twitter messages.* In "Proceedings of the first workshop on social media analytics". ACM.
21. Boettcher A and Lee D (2012). *EventRadar: A real-time local event detection scheme using twitter stream.* In "Green Computing and Communications (GreenCom), 2012 IEEE International Conference on", pp. 358-367. IEEE.
22. Rosi A, Mamei M, Zambonelli F, Dobson S, Stevenson G and Ye J (2011). *Social sensors and pervasive services: Approaches and perspectives.* In "Pervasive Computing and Communications Workshops (PERCOM Workshops), 2011 IEEE International Conference on", pp. 525-530. IEEE.
23. Colbaugh R and Glass K (2012). *Early warning analysis for social diffusion events.* "Security Informatics", Vol. 1(1).
24. Kwak H, Lee C, Park H and Moon S (2010). *What is Twitter, a Social Network or a News Media?* In "Proceedings of the 19th International Conference on World Wide Web", series WWW '10. ISBN 978-1-60558-799-8.
25. Brin, S., Page, L. (1998). *The anatomy of a large-scale hypertextual Web search engine.* Computer networks and ISDN systems, 30(1), 107-117.
26. Ye S and Wu SF (2010). *Measuring Message Propagation and Social Influence on Twitter.Com.* In "Proceedings of the Second International Conference on Social Informatics", series SocInfo'10. ISBN 3-642-16566-4, 978-3-642-16566-5.
27. Cha M, Haddadi H, Benevenuto F and Gummadi K (2010). *Measuring user influence in twitter: The million follower fallacy.* In "4th International AAAI Conference on Weblogs and Social Media (ICWSM)".
28. Romero DM, Meeder B and Kleinberg J (2011). *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". ISBN 978-1-4503-0632-4.
29. Wilson T, Wiebe J and Hoffmann P (2005). *Recognizing contextual polarity in phrase-level sentiment analysis.* In "Proceedings of the conference on human language technology and empirical methods in natural language processing", pp. 347-354. Association

for Computational Linguistics.
30. Lbner S (2000). *Polarity in natural language: predication, quantification and negation in particular and characterizing sentences.* "Linguistics and Philosophy", Vol. 23(3), pp. 213-308.
31. Pang B and Lee L (2008). *Opinion mining and sentiment analysis.* "Foundations and trends in information retrieval", Vol. 2(1-2), pp. 1-135.
32. Kouloumpis E, Wilson T and Moore J (2011). *Twitter sentiment analysis: The good the bad and the omg!* "ICWSM", Vol. 11, pp. 538-541.
33. Agarwal A, Xie B, Vovsha I, Rambow O and Passonneau R (2011). *Sentiment analysis of twitter data.* In "Proceedings of the Workshop on Languages in Social Media", pp. 30-38. Association for Computational Linguistics.
34. Baccianella S, Esuli A and Sebastiani F (2010). *SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining.* In "LREC", volume 10, pp. 2200-2204.
35. Miller GA (1995). *WordNet: a lexical database for English.* "Communications of the ACM", Vol. 38(11), pp. 39-41.
36. Yang QX, Yuan SS, Zhao L, Chun L and Peng S (2003). *Faster algorithm of string comparison.* "Pattern Analysis & Applications", Vol. 6(2).
37. Cavnar WB, Trenkle JM and others (1994). *N-gram-based text categorization.* "Ann Arbor MI", Vol. 48113(2).
38. Frnkranz J (1998). *A study using n-gram features for text categorization.* "Austrian Research Institute for Artifical Intelligence", Vol. 3(1998).

Fig. 5. Daily Average Severity Levels per Company and Airport

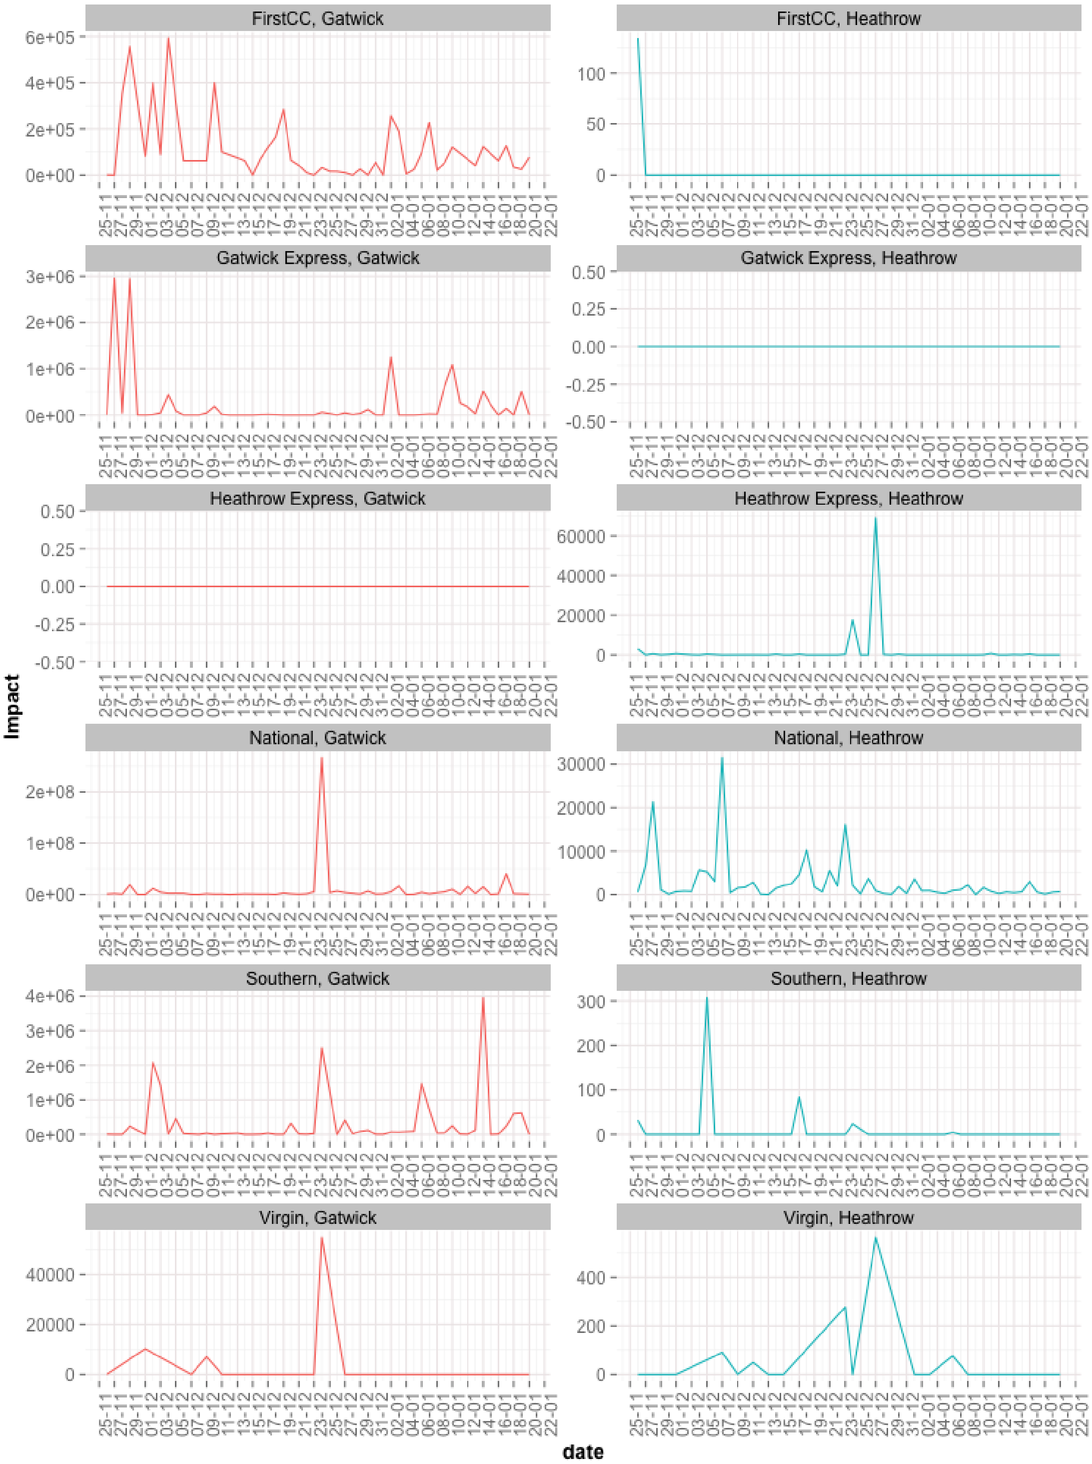

Fig. 6. Daily Impact per Company and per Airport

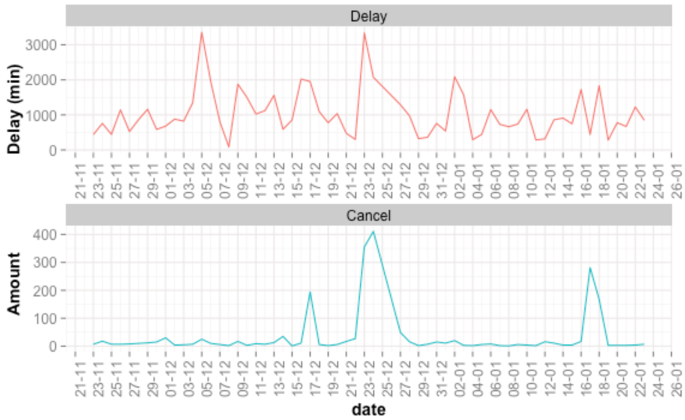

Fig. 7. Delay a) and Cancellations b) for First Capital Connect trains via Gatwick over 2 months. Source: http://www.firstcrapitalconnect.co.uk

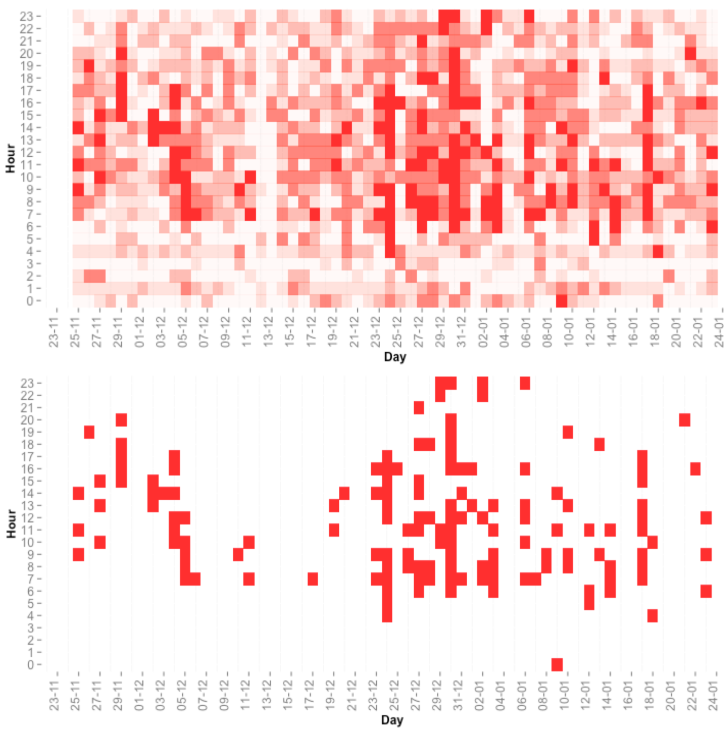

Fig. 8. Hourly heat-map based on 10 Impact Levels (a) and top decile (b) for National Railway in Gatwick

Bibliography

[ABL10] Ahn Y.-Y., Bagrow J. P., and Lehmann S. (2010) Link communities reveal multiscale complexity in networks. *Nature* 466(7307): 761–764.

[ABS+12] Aiello L. M., Barrat A., Schifanella R., Cattuto C., Markines B., and Menczer F. (2012) Friendship prediction and homophily in social media. *ACM Transactions on the Web (TWEB)* 6(2): 9.

[AH10] Asur S. and Huberman B. A. (2010) Predicting the future with social media. In *Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on*, volumen 1, pp. 492–499. IEEE.

[ASS12] Adam N. R., Shafiq B., and Staffin R. (2012) Spatial computing and social media in the context of disaster management. *IEEE Intelligent Systems* 27(6): 90–96.

[AXV+11] Agarwal A., Xie B., Vovsha I., Rambow O., and Passonneau R. (2011) Sentiment analysis of twitter data. In *Proceedings of the Workshop on Languages in Social Media*, pp. 30–38. Association for Computational Linguistics.

[BAM10] Bothos E., Apostolou D., and Mentzas G. (2010) Using social media to predict future events with agent-based markets. *IEEE Intelligent Systems* 25(6): 50–58.

[Bar69] Barnes J. A. (1969) Graph theory and social networks: A technical comment on connectedness and connectivity. *Sociology* 3(2): 215–232.

[Bas06] Basher R. (2006) Global early warning systems for natural hazards: systematic and people-centred. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 364(1845): 2167–2182.

[Ben12] Bennett S. (2012) Twitter now has more than 140 million active users sending 340 million tweets every day.

[BES10] Baccianella S., Esuli A., and Sebastiani F. (2010) Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *LREC*, volumen 10, pp. 2200–2204.

[BF06] Bussière M. and Fratzscher M. (2006) Towards a new early warning system of financial crises. *journal of International Money and Finance* 25(6): 953–973.

[BFK12] Bruce N. I., Foutz N. Z., and Kolsarici C. (2012) Dynamic effectiveness of advertising and word of mouth in sequential distribution of new products. *Journal of Marketing Research* 49(4): 469–486.

[BGK+05] Baumes J., Goldberg M. K., Krishnamoorthy M. S., Magdon-Ismail M., and Preston N. (2005) Finding communities by clustering a graph into overlapping subgraphs. *IADIS AC* 5: 97–104.

[BL99] Bradley M. and Lang P. (1999) Affective norms for english words (anew): Technical manual and affective ratings. *Gainsville, FL: Center for Research in Psychophysiology, University of Florida* .

[BM76] Bondy J. A. and Murty U. S. R. (1976) *Graph theory with applications*, volumen 290. Macmillan London.

[Bol98] Bolton R. N. (1998) A dynamic model of the duration of the customer's relationship with a continuous service provider: the role of satisfaction. *Marketing science* 17(1): 45–65.

[C+64] Coleman J. S. *et al.* (1964) Introduction to mathematical sociology. *London Free Press Glencoe.* .

[CCL10] Cheng Z., Caverlee J., and Lee K. (2010) You are where you tweet: a content-based approach to geo-locating twitter users. In *Proceedings of the 19th ACM international conference on Information and knowledge management*, pp. 759–768. ACM.

[CCMS10] Corley C. D., Cook D. J., Mikler A. R., and Singh K. P. (2010) Text and structural data mining of influenza mentions in web and social media. *International journal of environmental research and public health* 7(2): 596–615.

[CDH12] Candelon B., Dumitrescu E.-I., and Hurlin C. (2012) How to evaluate an early-warning system: Toward a unified statistical framework for assessing financial crises forecasting methods. *IMF Economic Review* 60(1): 75–113.

[Cen10] Centola D. (2010) The spread of behavior in an online social network experiment. *science* 329(5996): 1194–1197.

[CG12] Colbaugh R. and Glass K. (2012) Early warning analysis for social diffusion events. *Security Informatics* 1(1): 1–26.

[CHBG10] Cha M., Haddadi H., Benevenuto F., and Gummadi K. P. (2010) Measuring user influence in twitter: The million follower fallacy. In *ICWSM10: Proceedings of international AAAI Conference on Weblogs and Social*.

[CHLM10] Cavusoglu H., Hu N., Li Y., and Ma D. (Octubre 2010) Information technology diffusion with influentials, imitators, and opponents. *J. Manage. Inf. Syst.* 27(2): 305–334.

[CK02] Chen Y. and Kobayashi H. (2002) Signal strength based indoor geolocation. In *Communications, 2002. ICC 2002. IEEE International Conference on*, volumen 1, pp. 436–439. IEEE.

[CM06] Chevalier J. A. and Mayzlin D. (2006) The effect of word of mouth on sales: Online book reviews. *Journal of marketing research* 43(3): 345–354.

[CML11] Cho E., Myers S. A., and Leskovec J. (2011) Friendship and mobility: user movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1082–1090. ACM.

[CMZ10] Culnan M. J., McHugh P. J., and Zubillaga J. I. (2010) How large us companies can use twitter and other social media to gain business value. *MIS Quarterly Executive* 9(4): 112–115.

[CRA10] Clark E., Roberts T., and Araki K. (2010) Towards a pre-processing system for casual english annotated with linguistic and cultural information. In *Proceedings of the Fifth IASTED International Conference*, volumen 711, pp. 044–84.

[CSS+11] Chmiel A., Sobkowicz P., Sienkiewicz J., Paltoglou G., Buckley K., Thelwall M., and Hołyst J. A. (2011) Negative emotions boost user activity at bbc forum. *Physica A: statistical mechanics and its applications* 390(16): 2936–2944.

[Cul10] Culotta A. (2010) Towards detecting influenza epidemics by analyzing twitter messages. In *Proceedings of the first workshop on social media analytics*, pp. 115–122. ACM.

[CY14] Chan J. M. and Yazdanifard R. (2014) How social media marketing can influence the profitability of an online company from a consumer point of view. *Journal of Research in Marketing* 2(2): 157–160.

[Del03] Dellarocas C. (2003) The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management science* 49(10): 1407–1424.

[DHK+11] Dodds P. S., Harris K. D., Kloumann I. M., Bliss C. A., and Danforth C. M. (2011) Temporal patterns of happiness and information in a global social network: Hedonometrics and twitter. *PloS one* 6(12): e26752.

[EB99] Everett M. G. and Borgatti S. P. (1999) The centrality of groups and classes. *The Journal of mathematical sociology* 23(3): 181–201.

[FBPR10] Farris P. W., Bendle N. T., Pfeifer P. E., and Reibstein D. J. (2010) *Marketing metrics: The definitive guide to measuring marketing performance*. Pearson Education.

[Fle12] Fleurbaey M. (2012) The importance of what people care about. *Politics, Philosophy & Economics* 11(4): 415–447.

[FLG00] Flake G. W., Lawrence S., and Giles C. L. (2000) Efficient identification of web communities. In *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 150–160. ACM.

[FLGC02] Flake G. W., Lawrence S., Giles C. L., and Coetzee F. M. (2002) Self-organization and identification of web communities. *Computer* 35(3): 66–70.

[For10] Fortunato S. (2010) Community detection in graphs. *Physics Reports* 486(3): 75–174.

[FRG10] Flynn F. J., Reagans R. E., and Guillory L. (2010) Do you two know each other? transitivity, homophily, and the need for (network) closure. *Journal of personality and social psychology* 99(5): 855.

[FSB50] Festinger L., Schachter S., and Back K. (1950) The spatial ecology of group formation. *Social pressure in informal groups* pp. 33–60.

[GA05] Guimera R. and Amaral L. A. N. (2005) Functional cartography of complex metabolic networks. *Nature* 433(7028): 895–900.

[GAP$^+$10] Gobron S., Ahn J., Paltoglou G., Thelwall M., and Thalmann D. (2010) From sentence to emotion: a real-time three-dimensional graphics metaphor of emotions extracted from text. *The Visual Computer* 26(6-8): 505–519.

[GBG11] Gao H., Barbier G., and Goolsby R. (2011) Harnessing the crowdsourcing power of social media for disaster relief. *IEEE Intelligent Systems* 26(3): 10–14.

[GH12] Guille A. and Hacid H. (2012) A predictive model for the temporal dynamics of information diffusion in online social networks. In *Proceedings of the 21st International Conference Companion on World Wide Web*, pp. 1145–1152. ACM.

[GK09] Gilbert E. and Karahalios K. (2009) Predicting tie strength with social media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 211–220. ACM.

[GKMI$^+$10] Goldberg M., Kelley S., Magdon-Ismail M., Mertsalov K., and Wallace A. (2010) Finding overlapping communities in social networks. In *Social Computing (SocialCom), 2010 IEEE Second International Conference on*, pp. 104–113. IEEE.

[GKR00] Goldstein M., Kaminsky G. L., and Reinhart C. M. (2000) *Assessing financial vulnerability: an early warning system for emerging markets.* Peterson Institute.

[GMTA05] Guimera R., Mossa S., Turtschi A., and Amaral L. N. (2005) The worldwide air transportation network: Anomalous centrality, community structure, and cities' global roles. *Proceedings of the National Academy of Sciences* 102(22): 7794–7799.

[GN02] Girvan M. and Newman M. E. (2002) Community structure in social and biological networks. *Proceedings of the National Academy of Sciences* 99(12): 7821–7826.

[Gra73] Granovetter M. S. (1973) The strength of weak ties. *American journal of sociology* pp. 1360–1380.

[GRM$^+$12] Grabowicz P. A., Ramasco J. J., Moro E., Pujol J. M., and Eguiluz V. M. (2012) Social features of online networks: The strength of intermediary ties in online social media. *PloS (Public Library of Science) one* 7(1): 1–27.

[GW13] Gulati G. J. and Williams C. B. (2013) Social media and campaign 2012 developments and trends for facebook adoption. *Social Science Computer Review* 31(5): 577–588.

[GWDW09] Geuens M., Weijters B., and De Wulf K. (2009) A new measure of brand personality. *International Journal of Research in Marketing* 26(2): 97–107.

[Hay01] Haythornthwaite C. (2001) Tie strength and the impact of new media. In *Proceedings of the annual Hawaii International Conference on System Sciences*, page 11.

[HL71] Holland P. W. and Leinhardt S. (1971) Transitivity in structural models of small groups. *Comparative Group Studies* pp. 23–25.

[HN12] Hong S. and Nadler D. (2012) Which candidates do the public discuss online in an election campaign?: The use of social media by 2012 presidential candidates and its impact on candidate salience. *Government Information Quarterly* 29(4): 455–461.

[HPB13] Hays S., Page S. J., and Buhalis D. (2013) Social media as a destination marketing tool: its use by national tourism organisations. *Current Issues in Tourism* 16(3): 211–239.

[HTWF12] Hennig-Thurau T., Wiertz C., and Feldhaus F. (2012) 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* .

[Hub65] Hubbell C. H. (1965) An input-output approach to clique identification. *Sociometry* pp. 377–399.

[Hud08] Hudlicka E. (2008) What are we modeling when we model emotion? In *AAAI spring symposium: emotion, personality, and social behavior*, pp. 52–59.

[Hug05] Hughes A. (2005) *Strategic Database Marketing.* New York: McGraw-Hill.

[Jac01] Jaccard P. (1901) Étude comparative de la distribution florale dans une portion des alpes et des jura. *Bulletin del la Société Vaudoise des Sciences Naturelles* 37: 547–579.

[JLA11] Jin Y., Liu B. F., and Austin L. L. (2011) Examining the role of social media in effective crisis management: The effects of crisis origin, information form, and source on publics crisis responses. *Communication Research* page 0093650211423918.

[JLCD14] Jiang C., Liang K., Chen H., and Ding Y. (2014) Analyzing market performance via social media: a case study of a banking industry crisis. *Science China Information Sciences* 57(5): 1–18.

[JSFT07] Java A., Song X., Finin T., and Tseng B. (2007) 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*, pp. 56–65. ACM.

[JZSC09] Jansen B. J., Zhang M., Sobel K., and Chowdury A. (2009) Twitter power: Tweets as electronic word of mouth. *Journal of the American society for information science and technology* 60(11): 2169–2188.

[KA13] Kim Y. and Ahmad M. A. (2013) Trust, distrust and lack of confidence of users in online social media-sharing communities. *Knowledge-Based Systems* 37: 438–450.

[Kat53] Katz L. (1953) A new status index derived from sociometric analysis. *Psychometrika* 18(1): 39–43.

[KH10] Kaplan A. M. and Haenlein M. (2010) Users of the world, unite! the challenges and opportunities of social media. *Business horizons* 53(1): 59–68.

[KKT03] Kempe D., Kleinberg J., and Tardos É. (2003) Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 137–146. ACM.

[KLPM10] Kwak H., Lee C., Park H., and Moon S. (2010) What is twitter, a social network or a news media? In *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, pp. 591–600. ACM, New York, NY, USA.

[KRRT99] Kumar R., Raghavan P., Rajagopalan S., and Tomkins A. (1999) Trawling the web for emerging cyber-communities. *Computer networks* 31(11): 1481–1493.

[KS06] Kiss T. and Strunk J. (Diciembre 2006) Unsupervised multilingual sentence boundary detection. *Comput. Linguist.* 32(4): 485–525.

[Kum08] Kumar V. (2008) Customer lifetime value the path to profitability. *Foundations and Trends(R) in Marketing* 2(1): 1–96.

[KWM11] Kouloumpis E., Wilson T., and Moore J. (2011) Twitter sentiment analysis: The good the bad and the omg! *ICWSM* 11: 538–541.

[LBC98] Lang P. J., Bradley M. M., and Cuthbert B. N. (1998) Emotion, motivation, and anxiety: brain mechanisms and psychophysiology. *Biological psychiatry* 44(12): 1248–1263.

[LBdM13] Li T., Berens G., and de Maertelaere M. (2013) Corporate twitter channels: The impact of engagement and informedness on corporate reputation. *International Journal of Electronic Commerce* 18(2): 97–126.

[LH81] Leong T. and Holmest J. (1981) Communities of metazoan parasites in open water fishes of cold lake, alberta. *Journal of Fish Biology* 18(6): 693–713.

[Liu06] Liu Y. (2006) Word of mouth for movies: Its dynamics and impact on box office revenue. *Journal of marketing* 70(3): 74–89.

[LLN+95] Latané B., Liu J. H., Nowak A., Bonevento M., and Zheng L. (1995) Distance matters: Physical space and social impact. *Personality and Social Psychology Bulletin* 21(8): 795–805.

[LLVHB13] Leung D., Law R., Van Hoof H., and Buhalis D. (2013) Social media in tourism and hospitality: A literature review. *Journal of Travel & Tourism Marketing* 30(1-2): 3–22.

[Loe00] Loebner S. (2000) Polarity in natural language: predication, quantification and negation in particular and characterizing sentences. *Linguistics and Philosophy* 23(3): 213–308.

[MC84] Marsden P. V. and Campbell K. E. (1984) Measuring tie strength. *Social forces* 63(2): 482–501.

[McD11] McDonald M. (2011) *Marketing Accountability: A New Metrics Model to Measure Marketing Effectiveness.* Kogan Page Publishers.

[MD12] Meadows M. and Dibb S. (2012) Progress in customer relationship management adoption: a cross-sector study. *Journal of Strategic Marketing* 20(4): 323–344.

[Meh80] Mehrabian A. (1980) *Basic dimensions for a general psychological theory: Implications for personality, social, environmental, and developmental studies.* Oelgeschlager, Gunn & Hain, Incorporated.

[MF09] Mangold W. G. and Faulds D. J. (2009) Social media: The new hybrid element of the promotion mix. *Business horizons* 52(4): 357–365.

[Mil95] Miller G. A. (1995) Wordnet: a lexical database for english. *Communications of the ACM* 38(11): 39–41.

[MMG+07] Mislove A., Marcon M., Gummadi K. P., Druschel P., and Bhattacharjee B. (2007) Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pp. 29–42. ACM.

[MMM14] Middleton S., Middleton L., and Modafferi S. (2014) Real-time crisis mapping of natural disasters using social media. *IEEE Intelligent Systems* 99(PrePrints): (in press).

[MR74] Mehrabian A. and Russell J. A. (1974) *An approach to environmental psychology.* the MIT Press.

[MSLC01] McPherson M., Smith-Lovin L., and Cook J. M. (2001) Birds of a feather: Homophily in social networks. *Annual review of sociology* pp. 415–444.

[MT13] Marés J. and Torra V. (2013) On the protection of social networks users information. *Knowledge-Based Systems* 49: 134–144.

[MZ12] Macnamara J. and Zerfass A. (2012) Social media communication in organizations: The challenges of balancing openness, strategy, and management. *International Journal of Strategic Communication* 6(4): 287–308.

[New04] Newman M. E. (2004) Detecting community structure in networks. *The European Physical Journal B-Condensed Matter and Complex Systems* 38(2): 321–330.

[New06] Newman M. E. (2006) Modularity and community structure in networks. *Proceedings of the National Academy of Sciences* 103(23): 8577–8582.

[NGKA11] Naveed N., Gottron T., Kunegis J., and Alhadi A. C. (2011) Bad news travel fast: A content-based analysis of interestingness on twitter. In *Proceedings of the 3rd International Web Science Conference*, page 8. ACM.

[Nou13] Nousala S. (2013) Social adaptability of communities of practice case studies: A review of indicators and developments for economic sustainability. In *Proceedings of the 56th Annual Meeting of the ISSS-2012, San Jose, CA, USA*.

[OKA10] O'Connor B., Krieger M., and Ahn D. (2010) Tweetmotif: Exploratory search and topic summarization for twitter. In *ICWSM*.

[O'r07] O'reilly T. (2007) What is web 2.0: Design patterns and business models for the next generation of software. *Communications & strategies* (1): 17.

[PAMHV12] Porcel C., A. T.-L., Martnez M., and Herrera-Viedma E. (2012) A hybrid recommender system for the selective dissemination of research resources in a technology transfer office. *Information Science* 184(1): 1–19.

[PBMW98] Page L., Brin S., Motwani R., and Winograd T. (1998) The pagerank citation ranking: Bringing order to the web. In *Proceedings of the 7th International World Wide Web Conference*, pp. 161–172. Brisbane, Australia.

[PCKJ12] Park J., Cha M., Kim H., and Jeong J. (2012) Managing bad news in social media: A case study on domino's pizza crisis. In *ICWSM*.

[PKVS12] Papadopoulos S., Kompatsiaris Y., Vakali A., and Spyridonos P. (2012) Community detection in social media. *Data Mining and Knowledge Discovery* 24(3): 515–554.

[PL08] Pang B. and Lee L. (2008) Opinion mining and sentiment analysis. *Foundations and trends in information retrieval* 2(1-2): 1–135.

[PR05] Perry M. W. and Ricca A. V. (2005) Instant messaging: virtual propinquity for health promotion networking. *Promotion & education* 13(3): 211–212.

[PS12] Pan R. K. and Saramäki J. (2012) The strength of strong ties in scientific collaboration networks. *EPL (Euro Physics Letters)* 97(1): 18007.

[PT13] Paltoglou G. and Thelwall M. (2013) Seeing stars of valence and arousal in blog posts. *Affective Computing, IEEE Transactions on* 4(1): 116–123.

[RBDT12] Rajyalakshmi S., Bagchi A., Das S., and Tripathy R. M. (2012) Topic diffusion and emergence of virality in social networks. *CoRR* abs/1202.2215: 1212–2232.

[RBGH13] Rapp A., Beitelspacher L., Grewal D., and Hughes D. (2013) Understanding social media effects across seller, retailer, and consumer interactions. *Journal of the Academy of Marketing Science* 41(5): 547–566.

[RFPC07] Redondo J., Fraga I., Padrón I., and Comesaña M. (2007) The spanish adaptation of anew (affective norms for english words). *Behavior research methods* 39(3): 600–605.

[RGAH11] Romero D. M., Galuba W., Asur S., and Huberman B. A. (2011) Influence and passivity in social media. In *Proceedings of the 20th International Conference Companion on World Wide Web*, WWW '11, pp. 113–114. ACM, New York, NY, USA.

[RH11] Ramaswamy S. and Healey C. G. (2011) Visualization of the sentiment of tweets. In *SAS Analytics 2011 Conference*. SAS.

[RH14] Rust R. T. and Huang M.-H. (2014) The service revolution and the transformation of marketing science. *Marketing Science* 33(2): 206–221.

[RJ71] Richards Jr W. D. (1971) An improved conceptually-based method for analysis of communication network structure of large complex organizations. *International Communication Association* .

[RKJB13] 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): 108–127.

[RMK11] Romero D. M., Meeder B., and Kleinberg J. (2011) 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*, WWW '11, pp. 695–704. ACM, New York, NY, USA.

[RO15] Rana R. and Oliveira F. S. (2015) Dynamic pricing policies for interdependent perishable products or services using reinforcement learning. *Expert Systems with Applications* 42(1): 426–436.

[Ruh00] Ruhnau B. (2000) Eigenvector-centralitya node-centrality? *Social networks* 22(4): 357–365.

[Rus80] Russell J. A. (1980) A circumplex model of affect. *Journal of personality and social psychology* 39(6): 1161.

[Sab66] Sabidussi G. (1966) The centrality index of a graph. *Psychometrika* 31(4): 581–603.

[SBB+09] Scheffer M., Bascompte J., Brock W. A., Brovkin V., Carpenter S. R., Dakos V., Held H., Van Nes E. H., Rietkerk M., and Sugihara G. (2009) Early-warning signals for critical transitions. *Nature* 461(7260): 53–59.

[SBI+14] Sundsøy P., Bjelland J., Iqbal A. M., de Montjoye Y.-A., *et al.* (2014) Big data-driven marketing: How machine learning outperforms marketers gut-feeling. In *Social Computing, Behavioral-Cultural Modeling and Prediction*, pp. 367–374. Springer.

[SCC+13] Stefanidis A., Cotnoir A., Croitoru A., Crooks A., Rice M., and Radzikowski J. (2013) Demarcating new boundaries: mapping virtual polycentric communities through social media content. *Cartography and Geographic Information Science* 40(2): 116–129.

[SCCH09] Shen H., Cheng X., Cai K., and Hu M.-B. (2009) Detect overlapping and hierarchical community structure in networks. *Physica A: Statistical Mechanics and its Applications* 388(8): 1706–1712.

[SCP+12] Soares A. P., Comesaña M., Pinheiro A. P., Simões A., and Frade C. S. (2012) The adaptation of the affective norms for english words (anew) for european portuguese. *Behavior research methods* 44(1): 256–269.

[SDX13a] Stieglitz S. and Dang-Xuan L. (2013) Emotions and information diffusion in social mediasentiment of microblogs and sharing behavior. *Journal of Management Information Systems* 29(4): 217–248.

[SDX13b] Stieglitz S. and Dang-Xuan L. (2013) Emotions and information diffusion in social mediasentiment of microblogs and sharing behavior. *Journal of Management Information Systems* 29(4): 217–248.

[SGHVO+11] Serrano-Guerrero J., Herrera-Viedma E., Olivas J. A., Cerezo A., and Romero F. P. (2011) A google wave-based fuzzy recommender system to disseminate information in university digital libraries 2.0. *Information Sciences* 181(9): 1503–1516.

[SGLŽ14] Smailović J., Grčar M., Lavrač N., and Žnidaršič M. (2014) Stream-based active learning for sentiment analysis in the financial domain. *Information Sciences* page (in press).

[SL12] Shin H. and Lee J. (2012) Impact and degree of user sociability in social media. *Information Sciences* 196: 28–46.

[SNB+13] Silver D., Newnham L., Barker D., Weller S., and McFall J. (2013) Concurrent reinforcement learning from customer interactions. In *Proceedings of The 30th International Conference on Machine Learning*, pp. 924–932.

[SOM10a] Sakaki T., Okazaki M., and Matsuo Y. (2010) Earthquake shakes twitter users: real-time event detection by social sensors. In *Proceedings of the 19th international conference on World wide web*, pp. 851–860. ACM.

[SOM10b] Sakaki T., Okazaki M., and Matsuo Y. (2010) Earthquake shakes twitter users: Real-time event detection by social sensors. In *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, pp. 851–860. ACM, New York, NY, USA.

[SSL13] Stein A. D., Smith M. F., and Lancioni R. A. (2013) The development and diffusion of customer relationship management (crm) intelligence in business-to-business environments. *Industrial Marketing Management* 42(6): 855–861.

[SSP96] Smith W., Solow A. R., and Preston P. E. (1996) An estimator of species overlap using a modified beta-binomial model. *Biometrics* pp. 1472–1477.

[SSS74] Sutherland I. E., Sproull R. F., and Schumacker R. A. (Marzo 1974) A characterization of ten hidden-surface algorithms. *ACM Comput. Surv.* 6(1): 1–55.

[SV+04] Strapparava C., Valitutti A., *et al.* (2004) Wordnet affect: an affective extension of wordnet. In *LREC*, volumen 4, pp. 1083–1086.

[TDRM+06] Thomson M., Doblas-Reyes F., Mason S., Hagedorn R., Connor S., Phindela T., Morse A., and Palmer T. (2006) Malaria early warnings based on seasonal climate forecasts from multi-model ensembles. *Nature* 439(7076): 576–579.

[Tha89] Thayer R. E. (1989) *The biopsychology of mood and arousal.* Oxford University Press.

[Tic73] Tichy N. (1973) An analysis of clique formation and structure in organizations. *Administrative Science Quarterly* pp. 194–208.

[TLPP+14] Tejeda-Lorente A., Porcel C., Peis E., Sanz R., and Herrera-Viedma E. (2014) A quality based recommender system to disseminate information in a university digital library. *Information Science* 261: 52–69.

[VCK+09] Võ M. L., Conrad M., Kuchinke L., Urton K., Hofmann M. J., and Jacobs A. M. (2009) The berlin affective word list reloaded (bawl-r). *Behavior research methods* 41(2): 534–538.

[VYH08] Villanueva J., Yoo S., and Hanssens D. M. (2008) The impact of marketing-induced versus word-of-mouth customer acquisition on customer equity growth. *Journal of marketing Research* 45(1): 48–59.

[Was94] Wasserman S. (1994) *Social network analysis: Methods and applications*, volumen 8. Cambridge university press.

[Whi13] White D. (2013) Social media growth 2006-2012.

[WJ55] Weiss R. S. and Jacobson E. (1955) A method for the analysis of the structure of complex organizations. *American Sociological Review* 20(6): 661–668.

[WKB13] Warriner A. B., Kuperman V., and Brysbaert M. (2013) Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior research methods* 45(4): 1191–1207.

[WT85] Watson D. and Tellegen A. (1985) Toward a consensual structure of mood. *Psychological bulletin* 98(2): 219.

[WT02] Wu Y.-M. and Teng T.-l. (2002) A virtual subnetwork approach to earthquake early warning. *Bulletin of the Seismological Society of America* 92(5): 2008–2018.

[WWH05] Wilson T., Wiebe J., and Hoffmann P. (2005) Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the conference on human language technology and empirical methods in natural language processing*, pp. 347–354. Association for Computational Linguistics.

[WWL11] Wei W.-L., Wu C.-H., and Lin J.-C. (2011) A regression approach to affective rating of chinese words from anew. In *Affective Computing and Intelligent Interaction*, pp. 121–131. Springer.

[YLC+12] Yin J., Lampert A., Cameron M., Robinson B., and Power R. (2012) Using social media to enhance emergency situation awareness. *IEEE Intelligent Systems* 27(6): 52–59.

[YP11] Yates D. and Paquette S. (2011) Emergency knowledge management and social media technologies: A case study of the 2010 haitian earthquake. *International Journal of Information Management* 31(1): 6–13.

[YW10] Ye S. and Wu S. F. (2010) Measuring message propagation and social influence on twitter.com. In *Proceedings of the Second International Conference on Social Informatics*, SocInfo'10, pp. 216–231. Springer-Verlag, Berlin, Heidelberg.

[YYZ+03] Yang Q. X., Yuan S. S., Zhao L., Chun L., and Peng S. (2003) Faster algorithm of string comparison. *Pattern Analysis & Applications* 6(2): 122–133.

[ZFG11] Zhang X., Fuehres H., and Gloor P. A. (2011) Predicting stock market indicators through twitter "i hope it is not as bad as i fear". *Procedia-Social and Behavioral Sciences* 26: 55–62.