



Nowcasting earthquake damages with Twitter

Marcelo Mendoza^{1,3*} , Bárbara Poblete^{2,3} and Ignacio Valderrama²

*Correspondence:

marcelo.mendoza@usm.cl

¹Department of Informatics,
Universidad Técnica Federico Santa
María, Santiago, Chile

³Millennium Institute for
Foundational Research on Data
(IMFD), Santiago, Chile

Full list of author information is
available at the end of the article

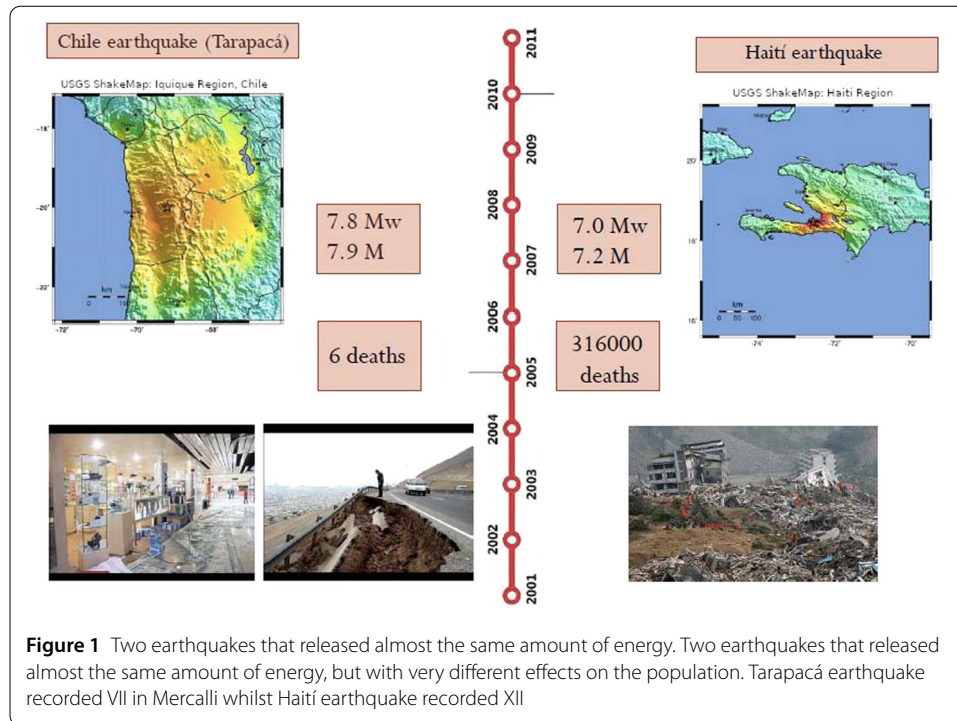
Abstract

The Modified Mercalli intensity scale (Mercalli scale for short) is a qualitative measure used to express the perceived intensity of an earthquake in terms of damages. Accurate intensity reports are vital to estimate the type of emergency response required for a particular earthquake. In addition, Mercalli scale reports are needed to estimate the possible consequences of strong earthquakes in the future, based on the effects of previous events. Emergency offices and seismological agencies worldwide are in charge of producing Mercalli scale reports for each affected location after an earthquake. However, this task relies heavily on human observers in the affected locations, who are not always available or accurate. Consequently, Mercalli scale reports may take up to hours or even days to be published after an earthquake. We address this problem by proposing a method for early prediction of spatial Mercalli scale reports based on people's reactions to earthquakes in social networks. By tracking users' comments about real-time earthquakes, we create a collection of Mercalli scale point estimates at municipality (i.e., state subdivisions) level granularity. We introduce the concept of *reinforced Mercalli support*, which combines Mercalli scale point estimates with locally supported data (named *'local support'*). We use this concept to provide Mercalli scale estimates for real-world events by providing smooth point estimates using a spatial smoother that incorporates the distribution of municipalities in each affected region. Our method is the first method based on social media that can provide spatial reports of damages in the Mercalli intensity scale. Experimental results show that our method is accurate and provides early spatial Mercalli reports 30 minutes after an earthquake. Furthermore, we show that our method performs well for earthquake spatial detection and maximum intensity prediction tasks. Our findings indicate that social media is a valuable source of spatial information for quickly estimating earthquake damages.

Keywords: Event damage assessment; Mercalli intensities

1 Introduction

The Modified Mercalli intensity scale ("Mercalli scale" for short) is an important measurement scale that summarizes the effects of an earthquake on public infrastructure, as well as in human damages. Unlike the moment magnitude scale, which quantifies the size of an earthquake in terms of released energy, the Mercalli scale is a qualitative measure that indicates perceived effects. Energy and damages do not always go hand-in-hand; an earthquake with the same magnitude in different regions may produce very different Mercalli scale measurements. This discrepancy is due to a number of physical variables,



such as the depth of the seismic movement and the geological properties of the soil for each location. In addition, damages to humans may vary, depending on construction standards used for buildings in each country. For instance, the 2005 Tarapacá earthquake occurred on June 13 with an epicenter located near locality of Mamiña, in the north of Chile, achieved 7.8 Mw in Richter (moment magnitude) producing 6 deaths. On the other hand, the 2010 Haiti earthquake occurred on January 24 with an epicenter located near Port-au-Prince achieved 7.0 Mw producing 316,000 deaths. Damages vary significantly between both events, as we show in Fig. 1.

Mercalli reports provide crucial information for timely emergency response and planning. Therefore, government agencies related to emergency management and geological centers strive to provide intensity reports in a timely and accurate manner to help mitigate disaster effects [1].

Mercalli reports are prepared by observers that have been appointed to different geographical areas. However, not all locations have appointed observers worldwide. Due to the human effort involved in producing intensity reports, these are commonly released hours or even days after a seismic movement. Many factors can obstruct the production of fast reports, among them the quality of communications during a disaster or the observer availability [2]. To improve intensity reporting, agencies such as the United States Geological Survey (USGS) have even created crowd-sourcing tools to collect this data online from regular people.^a On the other hand, social media users are regarded as providers of timely information, which allows the characterization of physical-world events [3]. Current advances show promising results in the direction of event analysis. Some of the most noteworthy studies address automatic event description and summarization using information provided by users as a situational information source [4]. However, social media data poses important challenges for information extraction, requiring researchers to design sophisticated methods for extracting useful knowledge from noisy data [5]. In the

specific scenario of earthquake disaster management, the state-of-the-art shows efforts in several tasks, such as earthquake detection and damage area detection [6–8], as well as maximum Mercalli intensity estimation [9]. These works are focused on the earthquake detection in a heavily populated city in which the earthquake was perceived by inferring the Mercalli intensity. As more populated cities may have many social media users, these events may produce a trending topic. Related work shows that it is possible to infer the maximum intensity in the Mercalli scale using trending topics. However, these methods have only been studied for the detection of high-energy events. Medium-scale events produce noisy data and often remain unexplored due to technical limitations.

We address the problem of early Mercalli scale estimation by studying how social media can contribute to this task. In particular, we propose a new approach, which focuses on the spatial estimation of Mercalli intensities and makes use of the volume and freshness of social data. We use municipalities as units of spatial aggregation, mapping posts to cities according to users' locations.

Social media data can be very noisy. We deal with noise by introducing the concept of “Reinforced Mercalli Support,” a key building block of our method. The idea behind “Reinforced Mercalli Support” is to process user posts that have high support at municipality level. Locally supported posts help us detect local trends, validating data that might be otherwise be ignored at more aggregated level of analysis (e.g., at state or country level). The spatial dimension of the problem, namely how people are distributed across a territory, and how this information affects the intensity inference process, is another key component of our proposal. We use spatial smoothing to deal with this aspect of the problem. Our method is inspired on the procedure used to elaborate Mercalli reports. These reports, which are based on data provided by on-the-scene experts, make use of the spatial distribution of the observers.

Our findings show that our approach can automatically provide spatial Mercalli reports in a timely and accurate manner. In particular, in this article we extend the following contribution introduced in our prior work [10]:

“We successfully deal with data at municipality level, detecting local trends in the specific task of maximum Mercalli intensity detection.”

The novel, previously unpublished, contributions of this current article are:

1. We present the first approach that addresses the problem of spatial Mercalli intensity inference based entirely on social media data.
2. We introduce a new concept, the *reinforced Mercalli support* estimate, which successfully combines local trends and local support in a single variable.
3. We show empirically, using real-world data, that our method provides accurate and fast Mercalli reports at a fine level of spatial granularity.
4. Our method can produce a spatial Mercalli report thirty minutes after an earthquake, contributing to improve current intensity estimation time and provides additional information to that given by human observers.

The paper is organized as follows. Section 2 presents a review of the relevant literature. In Sect. 3 we introduce our method proposal. In Sect. 4 we present our experimental validation, and we conclude in Sect. 5 with a discussion, conclusions, and outline of future work.

2 Related work

Twitter is an online social network with more than 300 million active users per month.^b This platform has become an huge source of real-time user-generated content. A sample of Twitter's streaming content is publicly available through the platform's API. These features have sparked notorious scientific interest during the past decade. Research includes, among others, studies to understand collective behavioral patterns [11], and to find correlations between social media and physical-world events [12].

During disaster situations and emergencies, many changes in the behavior of social media users have been observed, producing a mass convergence phenomena [3]. Social media has allowed researchers to gain insight into the dynamics of information propagation during crisis situations [5, 13, 14], displaying collaborative patterns useful for information filtering [15], and for the assessment of information credibility [16]. Social data has also been used to analyze, for example: forest fires [17], power outages in electrical systems [18], large-scale protests [19], and bus accidents in public transportation [20]. All of these systems place their efforts in the arena of providing situational awareness (local and timely information) of an event. Specifically, research has proven that social media can be valuable for rapidly assessing damage during large-scale disasters. Vieweg *et al.* [21] show how Twitter contributes enhance situational awareness during two natural hazards: the Oklahoma Grassfires and the Red River Floods in the U.S. Kryvasheyu *et al.* [22], on the other hand, studied Hurricane Sandy and discovered a strong relationship between hurricane-related Twitter activity and the actual path of the hurricane. Furthermore, they showed that for major disasters there is a correlation between damage and social media activity. Other efforts have focused on communication infrastructure during earthquakes by providing methods to favor message sharing during disasters in mobile networks [23], providing access to spatio-temporal data during disasters [24], and testing communication infrastructure using simulated data [25]. Along this line, several text mining methods have been used to elaborate reports, providing event summarization [4] (a short textual description of the earthquake) or detecting local related events such as looting and pillage [26].

Earthquake detection and analysis using social media is a particularly active field of study. The first efforts in this subject date from the year 2010, where the correlation between the number of tweets and the intensity of an earthquake was observed for the first time.^c In 2011, during the earthquake in Tohoku, researchers noted the existence of a high correlation between the number of user posts on Twitter (known as *tweets*) and the earthquake's intensity in certain locations [27, 28]. The relationship between tweet rates and Mercalli intensity was later revisited by Kropivnitskaya *et al.* [29]. They showed that tweets and Mercalli intensity correlated during three different earthquakes located in California, Japan, and Chile during 2014. An a related study, Crooks *et al.* [30] analyze the spatio-temporal characteristics of the relationship between an earthquake's seismic wave and social media posts. They show that Twitter data is comparable to that gathered by specialized crowdsourcing initiatives,^d and more rapid.

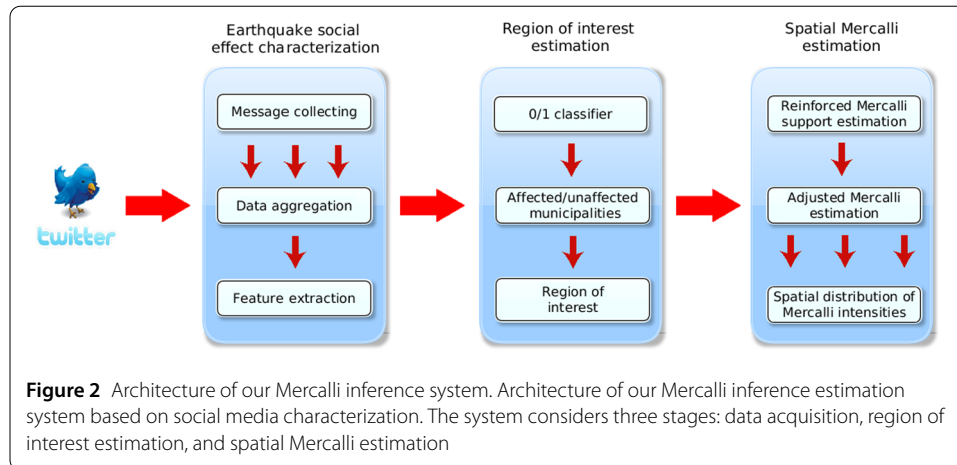
In addition, several earthquake alert systems have been created for different countries, such as for Australia [31], for Japan [6], and for Italy [7], as well as more general worldwide monitoring systems [8, 12]. Most of these systems use some type of *burst* detection algorithm over the tweet stream to report an earthquake, where a burst is defined as a large number of occurrences of tweets within a short time window [32]. These systems are fo-

cused on the specific task of earthquake detection, namely under which conditions we may confirm or deny that an earthquake reported in social media really occurred. Despite that the primary goal of these systems is to report that an earthquake happened in a given location, they have shown that it is possible to infer more information from social media data. Some of the most salient results on seismic event reports rely on the estimation of the epicenter of an earthquake using only information recovered from Twitter [6, 33]. Also, TwiFelt [34], which is an online system, uses the Twitter stream to estimate of the area in which an earthquake was felt in Italy. The system uses only geolocated tweets, with good performance for high-intensity earthquakes. However, reliability in this case depends on the existence of geolocated tweets, which can be scarce in many countries (between 4% and 7%) [35].

Regarding the problem of earthquake intensity estimation, Burks *et al.* [36] showed that Twitter can provide useful data to estimate shaking intensity. They proposed an approach that combines earthquake characteristics, measured using seismographs (such as *moment magnitude*, source-to-site distance, and wave velocity), with Twitter data (extracted from tweets that contained the term 'earthquake'). Conditioned by a set of reports retrieved from seismographs, they segmented the area around each recording station into nine radial subareas. They mapped to each of these areas, according to GPS location, all of the earthquake-related tweets produced during the 10-minutes posterior to an earthquake. They computed lexical features for each disc to study the correlation of these features with the Mercalli intensity. The authors showed good prediction of earthquake shaking intensity when combining earthquake measurements with tweets. In our work we use some of the same tweet features used by Burks *et al.* However, our method differs from theirs in that our goal is to provide rapid Mercalli estimates *only using social media data*, without seismograph recordings. We aim towards understanding the full contribution of social media for intensity estimation, as well as avoiding dependence on a dense seismographic networks.

Possibly, the closest to our proposal is the work of Cresci *et al.* [9]. In their approach, the authors studied how to estimate the maximum intensity in the modified Mercalli scale using only Twitter features. Using linear regression models over a collection of aggregated features, testing 45 different attributes. They showed that Twitter has enough predictive power to infer the maximum intensity of an earthquake. The set of features tested were extracted from user profiles, from tweet content, and from time-based features of the Twitter stream (e.g., tweet interval rates). Our proposal extends the work of by Cresci *et al.*, but with focus on the value of message content and on producing accurate spatially distributed estimations (not only maximum intensity prediction). Another difference is that our method uses only 12 lexical features, reducing dimensionality. In addition, we produce a spatial report of the event by enriching the process with spatial information like the geographical distribution of users. Nevertheless, in Sect. 4 we compare our method with Cresci *et al.* in the specific task of maximum intensity prediction.

In terms of spatially distributed data, we did not find at this moment prior work that generates spatial reports for earthquake intensity prediction. However, there is prior work that deal with producing spatial metrics, such as ours, for other types of natural hazards [22, 30, 35]. These works face similar methodological issues to our approach, related to accurate geolocation of messages, data density, and distance to the event location. Regarding the use of geolocation authors such as Yin *et al.* [37] have shown that location accuracy can



be improved by inferring locations from text at different geographical levels. For a more complete overview of the use of social media for mass emergencies and the challenges that this involves we suggest referring to the survey by Imran *et al.* [5]

3 Early inference of Mercalli intensities

We propose a methodology, based on social media user activity, to perform early inference of Mercalli intensities at municipality level. Social media data, in particular that of Twitter, can be rapidly collected and summarized. This allows us to produce Mercalli reports during the early stages of the aftermath of an earthquake. We divide our inference method into three stages, as is shown in Fig. 2: The first stage, *earthquake social effect characterization*, discussed in Sect. 3.1, is the process of collecting and aggregating Twitter data related to an earthquake occurrence. The data in this stage is aggregated at municipality level, which is the smallest geographical subdivision used by our approach. The second stage, *region of interest estimation*, discussed in Sect. 3.2, corresponds to the process of estimating which municipalities were actually affected by an earthquake. The third stage, *spatial Mercalli estimation*, discussed in Sect. 3.3, is in charge of predicting the spatial Mercalli estimates for the event. Next, we detail each of these stages.

3.1 Earthquake social effect characterization

Our approach uses Twitter as a data source of user-generated information about earthquakes. To obtain this data, every time an earthquake hits, we retrieve messages from the social platform that match any of the following keywords: *sismo*, *temblor*, *temblando* and *terremoto*, which loosely correspond to the terms *seismic*, *quake*, *shaking* and *earthquake* in Spanish. We collect these messages from the time of the earthquake up until 30 minutes afterwards. Given that our goal is to produce early Mercalli estimates we do not use any data posterior to 30 minutes after the event, because at that point the first partial Mercalli reports (produced by specialized agencies) start to appear.

Once the data has been collected, we aggregate it at *municipality* level, which is the smallest geographical subdivision used in our approach. We then process each municipality as an information unit, extracting features that describe how the earthquake affected that particular region. Table 1 details the features that we extract for each municipality for a particular earthquake.

Table 1 Municipality-level features per event. The first eleven features are calculated over the set of tweets that correspond to a specific municipality. The last feature corresponds to the municipality population

Feature	Description
NUMBER OF TWEETS	Nº of tweets produced in the municipality
TWEETS NORM	Nº of tweets produced in the municipality divided by the number of users in the municipality
AVERAGE WORDS	Avg. tweet length (in number of words)
AVERAGE LENGTH	Avg. tweet length (in number of chars)
QUESTION MARKS	<i>Fraction of tweets with...</i> ...question marks
EXCLAMATION MARKS	...exclamation marks
UPPER WORDS	...uppercase words
HASHTAG SYMBOLS	<i>Fraction of tweets containing the...</i> ...# (hashtag) symbol
MENTION SYMBOLS	...@ (mention) symbol
RT SYMBOLS	...RT symbol
WORD EARTHQUAKE	...word "earthquake"
POPULATION	Municipality population

In order to perform the municipality-level aggregation of data required by our approach we must examine each tweet for geolocation information. When available, the geolocation allows us to map a message back to the geographical area where it was originated. Hence, we will only use those messages for which we are able to extract a valid geolocation. In order to geolocate tweets we use the following steps: (1) if available, we extract the exact GPS coordinates from the tweet's *location* field, (2) if the location field was not provided by the user in their tweet, we then process the tweet's textual content. This is, we analyze the message's text (e.g., "*Earthquake in Valparaiso!!!*") to label possible location mentions using Named Entity Recognition (NER), then for each labeled location we use a fuzzy string matching procedure^e in order to map the location to its corresponding municipality. (3) At last, if all else fails, we apply the same procedure as in (2) but this time to the text provided by the user in their profile information. We acknowledge that this procedure can be noisy, since not all locations will be accurately mapped. However, we believe that spatial patterns will still emerge. In this sense, more accurate methods for tweet geolocation could improve this aspect of our approach. Nevertheless, this is an open problem that for the time being we consider beyond the scope of our work.

Once all of the remaining tweets have been aggregated at municipality level, each municipality is processed to extract 12 features, detailed in Table 1. These features provide a high-level characterization of user activity related to the earthquake, for each geographical subdivision.

3.2 Region of interest estimation

The next stage of our approach is estimating which municipalities were affected by the earthquake. We refer to these municipalities as the *region of interest* of an earthquake. Only those municipalities deemed as being affected by the earthquake will be used for spatial Mercalli intensity estimation in the following stage. To estimate the geographical subdivisions that were affected by the seismic event, we use a supervised classification model. This model separates municipalities into two classes: *unaffected* by the earthquake and *affected* by the earthquake.

To create this model we used a 0/1 classification algorithm, which we trained using municipality-level data modeled as feature vectors (using the features shown in Table 1).

The labels that we used for each municipality were class “0” if the earthquake was not perceived by the population (i.e., the municipality had no official Mercalli intensity value associated to it), and class “1” if the earthquake was perceived by the population (i.e., the municipality had an official Mercalli value associated to it). The Mercalli intensity values that we used to label the municipality-level data corresponded to values in official earthquake reports. More details on the technical and empirical aspects of the model creation are presented in Sect. 4.

3.3 Spatial Mercalli estimation

We next create spatial Mercalli estimates for the municipalities that are part of the region of interest of an earthquake. This process is divided into 3 steps: (i) *reinforced Mercalli support estimation*, (ii) *adjusted Mercalli estimation*, and (iii) *spatial distribution of Mercalli intensities*. We proceed to describe each of these steps.

3.3.1 Reinforced Mercalli support estimation

As a first step to estimate spatial Mercalli values, we define a municipality-level variable, which we call *reinforced Mercalli support*. The goal of this variable is to give more weight to intensity estimations that come from regions that displayed a larger amount of social activity. The rationale is to limit the effect of noisy reports by including only information with high local support.

Let i be the index that denotes a municipality belonging to the region of interest of a given earthquake. We then define the *local support* $s(i) \in [0, 1]$ of the i th municipality, as the ratio between users in i that reported the earthquake and the total number of different users in i which have reported earthquakes in the entire (training) dataset. Next, we define the *Mercalli point estimate* $m(i)$ for the i th municipality, as an intermediate estimate for the Mercalli value of i , which is obtained using a regression model. To estimate $m(i)$ we use a regression algorithm trained with earthquake Mercalli intensities and their corresponding municipality-level features. The Mercalli intensities used for each earthquake-municipality pair are based on official reports by governmental agencies. More details on the technical and empirical aspects of the regression model creation are discussed in Sect. 4.

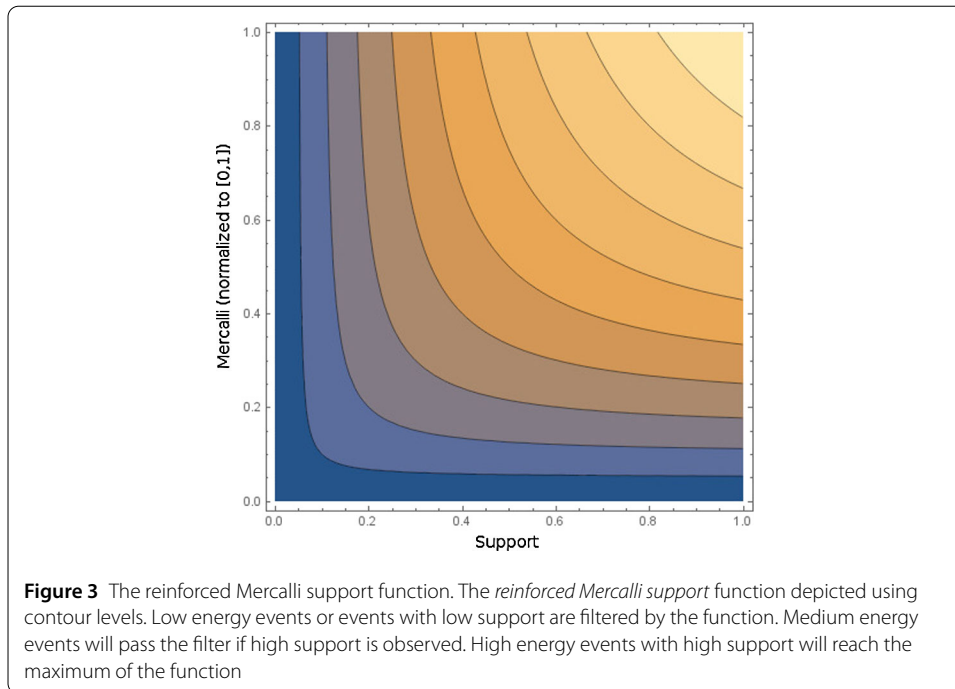
Next, we combine $s(i)$ and $m(i)$ to obtain the *reinforced Mercalli support* (m_{supp}) of i using a soft min function:

$$m_{\text{supp}}(i) = \frac{2 \cdot \bar{m}(i) \cdot s(i)}{\bar{m}(i) + s(i)}, \quad (1)$$

where $\bar{m}(i)$ is the Mercalli point estimate at the i th municipality bounded to the $[0, 1]$ interval. This is obtained by normalizing the Mercalli scale estimate from $\{1 \rightarrow 12\}$ to $[0, 1]$ using $\bar{m}(i) = \frac{m(i)-1}{11}$. Therefore, since $\bar{m}(i)$ and $s(i)$ are in the $[0, 1]$ range, the reinforced Mercalli support function is also in this range. We note that the reinforced Mercalli support function discards information from municipalities that do not have any local support (defined in Sect. 3.3). Figure 3 shows a contour level plot of this function.

3.3.2 Adjusted Mercalli estimation

Our method considers each municipality as an earthquake sensor. We model the activation of a sensor for a given earthquake using an activation function, in this case the sigmoid



function $\frac{1}{1+e^{-x}}$. We consider this sensor to be active when it reports a Mercalli intensity of 3 or greater, because 3 corresponds to the first level in the Mercalli scale for which an earthquake is perceived by humans. We incorporate this notion into our model by applying the sigmoid function to m_{supp} with an activation threshold of 3, i.e., $(11 \cdot m_{\text{supp}} + 1)_{(1,12)} - 2$ and combining this with the Mercalli point estimate $m(i)$. We refer to this value as the *adjusted Mercalli* estimate (m_{adj}) for municipality i :

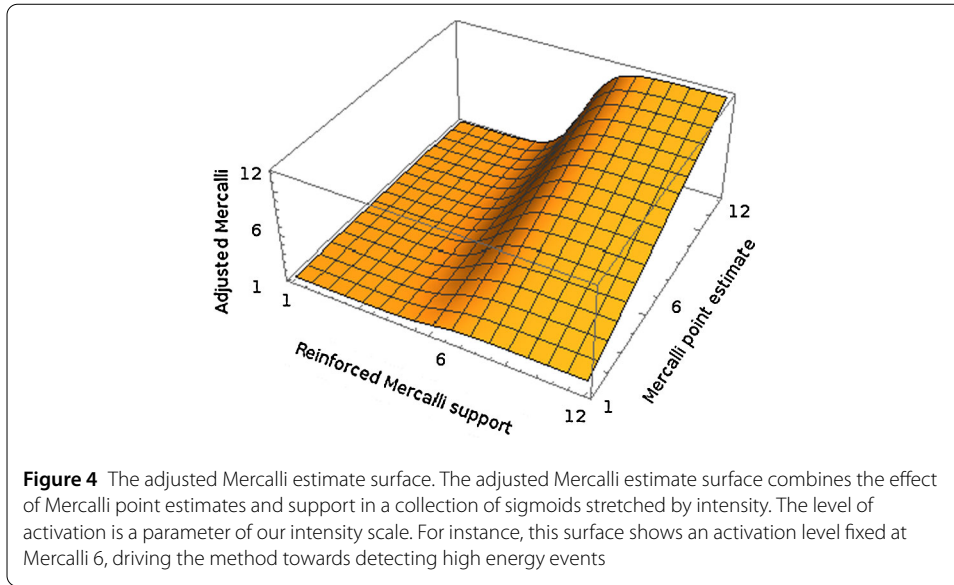
$$m_{\text{adj}}(i) = m(i) \cdot \text{Sigmoid}(11 \cdot [m_{\text{supp}}(i)] - 1). \quad (2)$$

The adjusted Mercalli value is obtained from a surface that comprises a collection of sigmoid functions in the Mercalli scale, stretching the sigmoid according to the Mercalli intensity, as we show in Fig. 4.

3.3.3 Spatial distribution of Mercalli intensities

Municipalities are defined as areal partitions of a geographical region. Hence, it makes sense that to predict municipality-level Mercalli intensities we consider the effect of its spatial correlations with other nearby municipalities that are part of the region of interest. To do this, we smooth the adjusted Mercalli estimate of each municipality in relation to the adjusted Mercalli estimates of its nearest neighbors. The influence of a neighbor on a given municipality is inversely determined by its geodesic distance to the municipality. This distance is measured as the pairwise distance between the largest city of the municipality and its neighbor. Since the adjusted Mercalli estimate (Eq. 2) of a municipality is conditioned to the local support that it had for the event (Eq. 2) by considering the largest city in each municipality, we give a high level of confidence to the distance estimation.

The idea behind this spatial smoothing is to provide a robust Mercalli estimation for municipalities that did not have sufficient support to provide a fair point estimate. This problem affects rural areas where Internet access is limited and/or marginal in proportion



to the municipality's population. In these locations, the use of spatial smoothing is helpful to infer a Mercalli intensity estimate even when a low number of Twitter reports are provided.

We describe the spatial smoothing process in detail next. First, we compute all possible pairwise geodesic distances between the municipalities. Then, for each municipality, we obtain its list of k -nearest neighbors, denoted as $k\text{-nn}(i)$, where k is a parameter. For each municipality i , we normalize its distance to each of its k -nearest neighbors by the sum of the distances to all of the $k\text{-nn}(i)$:

$$d(i, j) = \frac{d_{\text{geo}}(i, j)}{\sum_{j' \in k\text{-nn}(i)} d_{\text{geo}}(i, j')}. \quad (3)$$

Note that $\sum_{j \in k\text{-nn}(i)} d(i, j) = 1$. To model the influence of each neighbor on the municipality, we convert the distance into a similarity, as follows:

$$\text{sim}(i, j) = \frac{1 - d(i, j)}{\sum_{j' \in k\text{-nn}(i)} 1 - d(i, j')}. \quad (4)$$

Finally, we create a smoothed Mercalli (m_{sm}) point estimate for i using a linear combination of the adjusted Mercalli point estimate of i and of its neighbors:

$$m_{\text{sm}}(i) = (1 - \lambda) \cdot m_{\text{adj}}(i) + \lambda \cdot \sum_{j' \in k\text{-nn}(i)} \text{sim}(i, j') \cdot m_{\text{adj}}(j'), \quad (5)$$

where λ is a parameter within $[0, 1]$ that controls the relative weight given to the point estimate and the neighborhood. As the Mercalli scale takes integers values, we round $m_{\text{sm}}(i)$ to its closest integer in the modified Mercalli scale $\{1 \rightarrow 12\}$.

4 Experiments

In this section we present the experimental validation of our proposed method. In particular, we evaluate the performance of our approach for estimating Mercalli values at

municipality level based on social media activity during earthquakes. First, we present a description and characterization of our datasets, and secondly we present our results.

4.1 Dataset description and characterization

We use two datasets, the first a ground truth dataset obtained from a seismological agency and the second is a Twitter dataset from which our method performs its Mercalli estimations. We describe them next.

4.1.1 Ground truth earthquake dataset

As a ground truth dataset we used an earthquake catalog provided publicly by the National Seismological Center of Chile, also known internationally as GUC. This catalog contains information about earthquakes registered in Chile from January 2016 to June 2017. This information is provided at municipality level and includes event magnitude, reported in Moment Magnitude scale and, if the earthquake was perceived by the population, it contains as well its Mercalli intensity report. The catalog contains 332 earthquakes perceived by the population, which ranged from 2.2 M_l to 7.6 M_w in magnitude. Each entry in the catalog corresponds to a earthquake-municipality pair with its corresponding intensity value in the Mercalli scale. In total, the catalog comprises 8296 entries. We use a local-scope catalog because this catalog contains fine-grained data about earthquakes in the Chilean territory for all magnitude ranges, which are not otherwise available in global-scope catalogs.

4.1.2 Twitter dataset

Our second dataset corresponds to data obtained from the public Twitter stream, using the search API.^f In order to retrieve conversations related to earthquakes, we collected tweets that matched any of the following keywords (in Spanish) *seismic*, *quake*, *shaking* and *earthquake*. Overall, we collected 825,310 tweets, which were posted by 309,749 different users during the time period of our study (i.e., from January 2016 to June 2017). From these tweets, we wanted to keep only those that corresponded to earthquake mentions generated in Chile, so we could use them with our local-scope ground truth data. However, only 2200 of these tweets had GPS locations (0.26%). Therefore, we extracted additional location information from users' profiles using the heuristic approach described in detail in Sect. 3.1. Using this approach we found that 207,015 users (i.e., 66.8%) registered a valid location in their profile, of which 57,546 indicated to belong to Chile. For the users that indicated being in Chile, we then used approximate matching to associate their profile information to a list of Chilean municipalities. This resulted in a total match of 41,885 users to Chilean municipalities, which in turn yielded a total of 187,317 tweets mapped to 345 different municipalities in Chile.

Next, we performed a match between the earthquakes in our ground truth earthquake dataset and our Twitter dataset. We matched each entry in the ground truth catalog to its corresponding Twitter data municipality, to create municipality-level Twitter data units, as described in Sect. 3. To create municipality-level Twitter data units we considered tweets from the time at which an earthquake occurred until 30 minutes afterwards. This is, each information unit is composed by no more than 30 minutes of tweets. Overall, this process resulted in a total of 6790 municipality-earthquake pairs with Mercalli information, and 6548 municipality-earthquake pairs without Mercalli information (i.e.,

Table 2 Our dataset in terms of municipality-level local instances and the coverage of Twitter over the GUC catalog

Intensity	1	2	3	4	5	6	7	Overall
GUC	2031	2212	2430	1121	355	133	14	8296
Twitter + GUC	1033	1808	2334	1116	352	133	14	6790
Not Covered	998	404	96	5	3	–	–	1506
Coverage (%)	50.8	81.7	96	99.5	99.1	99.1	100	81.8

6790/6548 *affected/unaffected* municipality information units). The number of data units of our dataset is shown in Table 2 disaggregated per Mercalli intensity level.

Table 2 shows number of data units at municipality level for the GUC catalog (“GUC”), the intersections of both datasets (“Twitter + GUC”), the number of GUC events that did not have coverage on Twitter (“Not Covered”), and the percentage of units covered (“Coverage %”). According to Table 2, our dataset has large coverage of the events registered by the GUC, with almost perfect coverage for medium to high-energy events. Low-energy events are less reported on Twitter because many times they are not perceived by the population. Overall, the intersection between Twitter and the GUC data produces 6790 municipality-level data instances, with an average coverage of 81.8%.

4.1.3 Data characterization

We first performed a data exploration process to analyze the relationship between municipality-level features and Mercalli values. We studied the existence of correlations between municipality-level propagation features, shown in Table 3.

The first 2-row block in Table 3 shows the correlation between the first variable, the target variable `MERCALLI`, and each of the twelve features used by our method. The second 2-row block shows correlation between the second variable, `NUMBER OF TWEETS`, and each of the remaining features (except `MERCALLI`), and so on for the rest of the table.

There is a positive correlation between `NUMBER OF TWEETS` and `TWEETS NORM`. There is a negative correlation between `TWEETS NORM` and `POPULATION`. An expected correlation arises between `AVERAGE WORDS` and `AVERAGE LENGHT`, and also between `MENTION SYMBOLS` and `RT SYMBOLS`, since the latter is a subset of the former. This is because messages that are re-posted on Twitter (i.e., *retweeted*) always include a mention to the author of the original message. As expected, the number of special symbols increases with tweet length.

In Fig. 5, we show boxplots for each feature in relation to Mercalli intensity values. First, we discuss the median, which as expected, for `NUMBER OF TWEETS` and `TWEETS NORM` is correlated with `MERCALLI`. The use of normalization for `TWEETS NORM` reduced the observed variance, shown in the relatively similar sizes of each box for this feature. The medians of the `AVERAGE LENGTH` and `AVERAGE WORDS` plots tend to decrease when `MERCALLI` increases. We say that this is only a tendency because high-energy events do not follow that pattern. In particular, for earthquakes level 7 in Mercalli, the median increases. The inversion of the pattern for high-energy earthquakes is also observed in other features (see, for instance, `NUMBER OF TWEETS` and `EXCLAMATION MARKS`). In the plots for `QUESTION MARKS` and `EXCLAMATION MARKS`, the median increases during high-energy earthquakes. The use of uppercase is marginal in the dataset, showing more presence during high-energy earthquakes. A similar pattern is observed in the boxplots

Table 3 Spearman correlation coefficient for the features considered in our study. As expected, the correlation between NUMBER OF TWEETS and TWEETS NORM is strong, as well as that between AVERAGE WORDS and AVERAGE LENGHT, and between MENTION SYMBOLS and RT SYMBOLS. The Spearman coefficients found are statistically significant, as the p -values show. Strong correlations are indicated with bold fonts

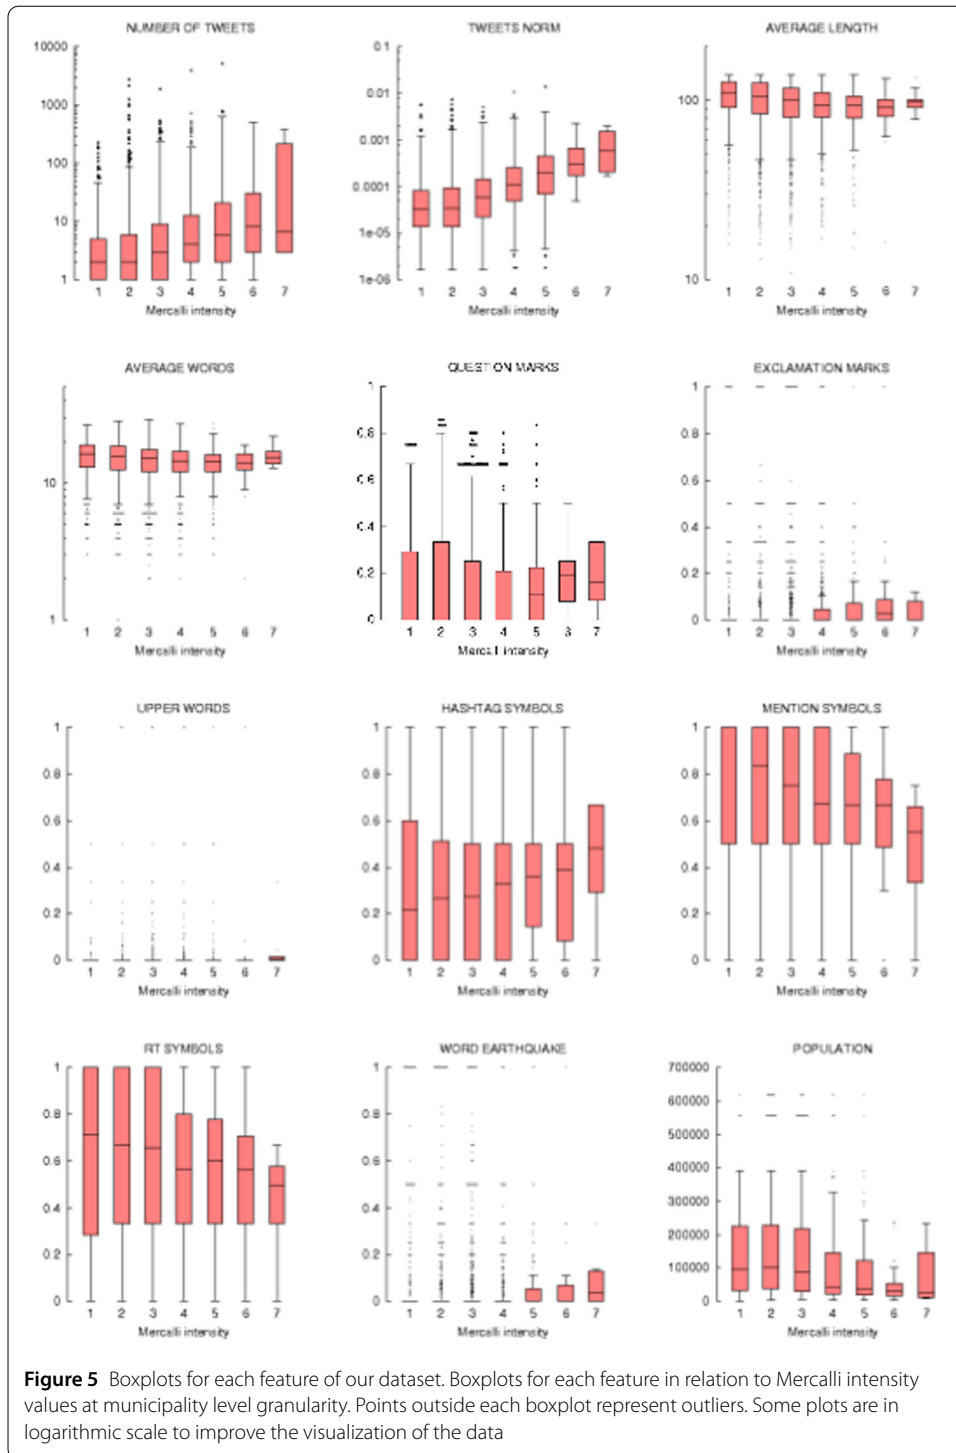
	NUMBER OF TWEETS MERCALLI	TWEETS NORM	AVERAGE WORDS	AVERAGE LENGHT	QUESTION MARKS	EXCLAMATION MARKS	UPPER WORDS	HASHTAG SYMBOLS	MENTION SYMBOLS	RT SYMBOLS	WORD EARTHQUAKE	POPULATION
ρ	0.22	0.24	-0.11	-0.14	0.05	0.16	0.15	0.06	-0.12	-0.11	0.05	-0.02
p	2.2e-16	2.2e-16	2.2e-16	2.2e-16	3.3e-13	2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16	6.2e-12	0.0
ρ		0.52	-0.09	-0.18	0.29	0.38	0.22	0.21	-0.22	-0.17	0.25	0.3
p		2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16
ρ			-0.07	-0.16	0.19	0.29	0.18	0.12	-0.14	-0.11	0.14	-0.56
p			2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16
ρ				0.86	0.02	-0.11	-0.08	-0.04	0.31	0.33	0.02	0.0
p				2.2e-16	9.5e-05	2.2e-16	2.2e-16	3.1e-10	2.2e-16	2.2e-16	0.002	0.201
ρ					-0.08	-0.18	-0.1	-0.04	0.41	0.43	0.0	0.0
p					2.2e-16	2.2e-16	2.2e-16	1.1e-09	2.2e-16	2.2e-16	0.249	0.956
ρ						0.23	0.08	0.13	-0.04	-0.01	0.1	0.06
p						2.2e-16	2.2e-16	2.2e-16	3.1e-10	0.031	2.2e-16	2.2e-16
ρ							0.18	0.11	-0.14	-0.12	0.23	0.09
p							2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16	2.2e-16
ρ								0.04	-0.07	-0.05	0.12	0.09
p								6.7e-10	2.2e-16	4.1e-11	2.2e-16	2.2e-16
ρ									-0.13	-0.07	0.07	0.06
p									2.2e-16	2.2e-16	2.2e-16	2.2e-16
ρ										0.86	-0.07	-0.07
p										2.2e-16	2.2e-16	2.2e-16
ρ											-0.07	-0.04
p											2.2e-16	9.8e-09
ρ												0.15
p												2.2e-16

for WORD EARTHQUAKE, which increases when MERCALLI increases. The use of HASH-TAG SYMBOLS increases with MERCALLI, as expected. However, for MENTION SYMBOLS and RT SYMBOLS the median decreases with MERCALLI. Finally, POPULATION shows a clear inverse relation of the median and MERCALLI, reinforcing the presence of a negative linear correlation, as was shown in Table 3.

A second aspect that we analyzed was the variance. Boxplots in Fig. 5 show low variance with MERCALLI in several cases (see boxplots for TWEETS NORM, AVERAGE LENGHT, and AVERAGE WORDS). However, there are features which show high variance in relation to MERCALLI, such as HASHTAG SYMBOLS, RT SYMBOLS, and MENTION SYMBOLS.

4.2 Experiment and results

From the total of 332 earthquakes, 264 were selected for training tasks, and the remaining 68 earthquakes were reserved for testing and validation tasks. This represents a training/testing split of 80/20 percent. The training/testing split process was conducted using stratified random sampling over earthquakes according to each Mercalli intensity level, keeping the same proportions between intensities in training and testing folds, avoid-



ing over/under representations of low/high energy earthquakes in training and/or testing folds. Training/testing proportions of instances according to the maximum Mercalli intensity of each earthquake are shown in Table 4.

Table 4 Training/testing partition for events according to the maximum Mercalli intensity of each seismic movement. High energy events are less frequent than low energy events. Note that this table summarizes our dataset in terms of number of earthquakes, but for each of these earthquakes we have many municipality-level local instances. In fact, as high energy earthquakes cover a wider area, they produce many municipality-level instances

Max intensity	2	3	4	5	6	7
Training	11	105	103	39	4	2
Testing	3	26	26	10	2	1
Overall	14	131	129	49	6	3

Table 5 Training accuracy per class for the *region of interest* estimation, using 5-fold cross validation

Class	FP Rate	Precision	Recall	F-measure	ROC Area
0	0.189	0.736	0.575	0.646	0.693
1	0.425	0.675	0.811	0.737	0.693
W. Avg.	0.312	0.705	0.698	0.693	0.693

Table 6 Testing accuracy per class for the region of interest detection task. Low precision indicates that the problem is hard to solve using classification at fine level granularity. However, a simple 0/1 classifier is sufficient to infer the region of interest with **0.816** recall

Class	FP Rate	Precision	Recall	F-measure	ROC Area
0	0.184	0.765	0.517	0.617	0.667
1	0.483	0.594	0.816	0.687	0.667
W. Avg.	0.323	0.685	0.656	0.650	0.667

4.2.1 Region of interest estimation

Training/testing municipality data batches accounted for 10,491/2847 instances at municipality level. As detailed in Sect. 3.2, in order to define the *region of interest* of an earthquake, we train a 0/1 classifier to estimate which municipalities were affected/unaffected by earthquakes. In the training fold, 5021 instances corresponded to class 0 (unaffected) and 5470 to class 1 (affected). Using Weka 3.7,⁸ we trained a support vector machine (SVM) using five-fold cross validation, using C-SVC (C-support vector classification) with a radial basis function kernel. Since our focus was to detect class 1 instances, we used cost-sensitive learning, penalizing class 1 false negatives to maximize recall, at the cost of an increased FP-rate (false positive rate). We tested other algorithms such as naive Bayes and multilayer perceptron; SVM displayed the best results, with 7325 correctly classified instances, representing an overall accuracy of 69.82%. Table 5 shows detailed accuracy per class.

We applied the resulting model, obtaining 1867 correctly classified instances, over a total of 2847 instances, achieving an accuracy of 65.57%. This shows that the classifier generalizes well, since the overall accuracy of training and testing partitions are similar. Low precision for this task illustrates that this problem is difficult to solve probably due to the presence of noise at fine level aggregation. However, what remains important is that the testing recall is high, indicating good predictability for class 1.

Nevertheless, Table 6 shows the testing accuracy per class, indicating that the simple 0/1 classifier is enough to recover the region of interest of an earthquake with good recall (0.816). Hence, each region of interest is over-estimated (shown by the low precision on class 1), but it still achieves good coverage of the actual region of interest (shown by the high recall).

Table 7 Testing performance according to the actual Mercalli intensity for the region of interest detection task. A 0/1 classifier is sufficient for medium and high energy events at municipality level, showing good performance in terms of recall

Act.	Pred.	-	1	2	3	4	5	6	7
0	0	790	–	–	–	–	–	–	–
0	1	737	–	–	–	–	–	–	–
1	0	–	66	85	62	25	5	–	–
1	1	–	130	234	351	198	65	95	4
Instances		1527	196	319	413	223	70	95	4
Error rate		0.48	0.33	0.26	0.15	0.11	0.07	–	–

To better understand how the 0/1 classifier behaves, in Table 7 we disaggregated testing instances according to the actual level of Mercalli intensity. We can observe that false negative rate is very low, and as the intensity of the earthquake increases, the error rate decreases. High-energy earthquakes (from Mercalli 5 and above) present almost perfect performance. The largest part of the error occurs for low-intensity earthquakes (from Mercalli 3 and below). This is intuitive since for this segment of the Mercalli scale most people will not perceive the earthquake, because it can only be perceived under extremely favorable conditions (e.g., in the top floor of a building). Conversely, for the 0/1 classifier, it is difficult to distinguish municipalities that reported the earthquake in their social networks, but in which the actual event was not perceived, producing false positives. We handle this overestimation of the region of interest using spatial smoothing, described next.

4.2.2 Using regression and spatial smoothing to estimate Mercalli intensities

As detailed in Sect. 3.3, we performed the experimental validation for regression and spatial smoothing procedures to estimate Mercalli intensities. We used a support vector regression model with a sequential minimal optimization (SMO) algorithm implemented in Weka 3.7. We trained using five-fold cross-validation using as training instances municipalities where an earthquake was perceived with Mercalli values ranging from 1 to 7. To deal with intensity unbalance, we applied instance re-sampling biased to class uniformity, achieving a total of 5470 training instances. To calculate the support vectors, we used a normalized polynomial kernel with an exponent equal to 2. During the training process, the fitted model achieved a correlation coefficient of 0.65 with mean absolute error (MAE) of 1.15. The same configuration was used to fit an SMO regression model over a reduced set of features, with high correlation with the Mercalli intensity (NUM TWEETS, NUM TWEETS NORM, and POPULATION), achieving a correlation coefficient of only 0.304. Therefore, after corroborating that the best values for correlation coefficient in the regression were achieved using all features, we discarded the model based on the reduced set of features. We selected the model based on all 12 features as a baseline predictor of Mercalli at municipality level (denoted by $m(i)$ in Equation (2)).

After re-evaluation on the test set, the correlation coefficient decreased to 0.26 with a MAE of 2.26. This result indicates that the sole use of a regression is insufficient to perform accurate predictions. We show next that the use of our adjusted Mercalli estimation and the inclusion of spatial smoothing boosts the method's accuracy, outperforming the baseline.

After calculating the reinforced Mercalli estimate and the adjusted Mercalli, we applied spatial smoothing using k -NN with $k = 5$. Then, we tuned λ by evaluating the MAE measured between the actual Mercalli and the Mercalli estimate given by our method. For

Table 8 Overall MAE at different values of λ

λ	0	0.2	0.4	0.6	0.8	1.0
Overall MAE	2.078	1.841	1.551	1.207	0.876	1.029

each earthquake in the testing set, we averaged the MAE over municipalities and obtained a single MAE estimate per earthquake. Then, we stratified the MAEs via the maximum Mercalli intensity per earthquake and calculated the MAE estimation per Mercalli level. (For instance, MAE(2) is the MAE averaged over all the earthquakes in the testing set with maximum intensity 2 in the Mercalli scale). As the distribution of earthquakes per Mercalli level of intensity is unbalanced, we weighted the error in proportion to the intensity of the earthquake, paying high costs in high-intensity earthquakes and low costs in low-intensity earthquakes. We named this measure OVERALL MAE and defined it as follows:

$$\text{OVERALL MAE} = \frac{\sum_{M \in \text{MERCALLI SCALE}} \text{MAE}(M) \cdot M \cdot \#\text{INSTANCES AT } M}{\sum_{M \in \text{MERCALLI SCALE}} M \cdot \#\text{INSTANCES AT } M}.$$

The OVERALL MAE for values of λ ranging in $\{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$ is shown in Table 8.

Table 8 shows the value of using spatial smoothing for our method. On the one hand, when spatial smoothing is dismissed ($\lambda = 0$), the method achieves its worst result with an OVERALL MAE at 2.078, which is almost the same value achieved using the baseline. On the other hand, when the prevalence of the spatial component increases, the error decreases. The best value is achieved at $\lambda = 0.8$.

Now, we show the disaggregated error at each level of Mercalli intensity. For each earthquake in the testing set of earthquakes, we measured the absolute error in each municipality of Chile. These results are shown in Fig. 6. Figure 6 shows the errors stratified per intensity, where each boxplot indicates errors related to earthquakes with the intensity indicated in the x-axis. The figure shows that our proposal outperforms the baseline in all the comparisons, which is only based on regression at municipality level. For instance, for intensity 5, our method shows its best performance in an earthquake with an absolute error of 0.2 (median minus deviation). The median of the boxplot is located at 1 meaning that in average our method records an absolute error of one degree at this level of the Mercalli intensity scale. This result indicates that the use of highly supported regressors helps to reduce the estimation error. The prediction of medium-scale energy events shows great performance (see the error measures achieved in the range of $\{3, 5\}$), outperforming the baseline by up to 2 points of error, on average.

To better understand the quality of our results, note that the results displayed in Fig. 6 correspond to the errors of a complete spatial Mercalli report, which is a result significantly distinct from those observed in state-of-the-art methods, where the Mercalli estimation is focused on the aggregated estimation of the maximum intensity per earthquake. We remark at this point that our method is the first to provide a complete spatial Mercalli report.

To verify the existence of sufficient evidence to sustain that our method outperforms the baseline, we conducted a non parametrical statistical test of differences in absolute errors. To do this we conducted a Wilcoxon rank sum tests on errors stratified per intensity. Wilcoxon statistics, p -values and 95 percent confidence intervals of location displacement at each level of intensity are shown in Table 9.

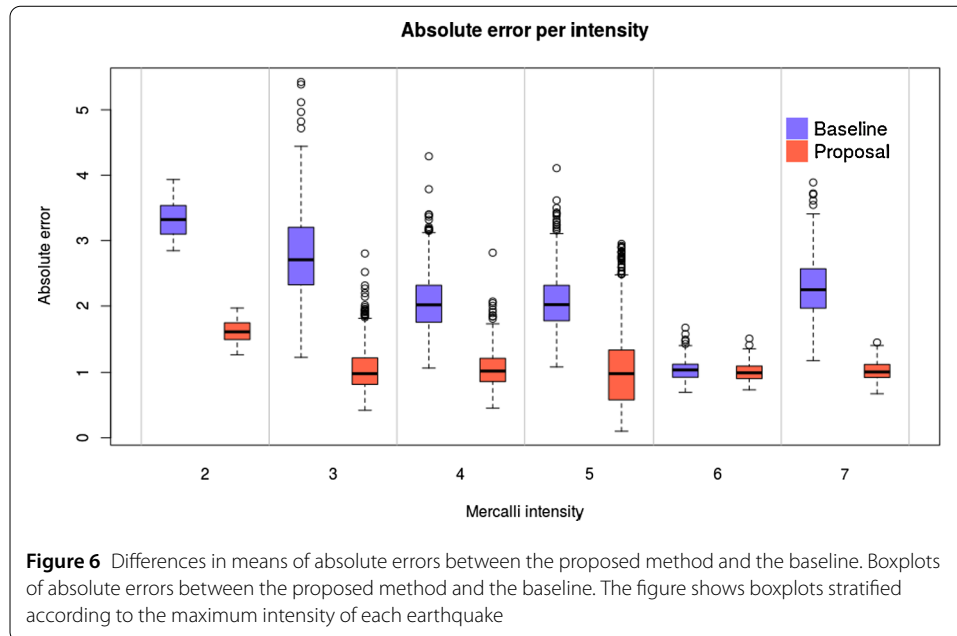


Table 9 Wilcoxon rank sum test results for differences in absolute errors between our method and the baseline. The second column indicates the Wilcoxon statistic per each level of the test

	W	p-value	Displacement	Confidence interval
2	0	2.2e-16	-1.68	[-1.81, -1.57]
3	1987	2.2e-16	-1.71	[-1.77, -1.63]
4	13,709	2.2e-16	-0.99	[-1.02, -0.96]
5	149,490	2.2e-16	-1.02	[-1.07, -0.96]
6	22,640	0.03692	-0.02	[-0.05, -0.01]
7	61	2.2e-16	-1.22	[-1.28, -1.16]

4.2.3 Maximum intensity prediction task using social media

We compare our method with the state-of-the-art technique for the specific task of maximum intensity prediction using social media. According to our literature review, the method proposed by Cresci *et al.* [9] is the only one that we can be directly compared to ours, in the sense that it relies *solely* on Twitter data; as such, it can be evaluated using the same dataset used by our method. As we noted in our related work overview, in Sect. 2, the work by Burks *et al.* [36] also uses Twitter data to estimate intensity. However, their approach is not directly comparable to ours because their model uses Twitter in combination with actual seismograph measurements. Since the goal of our work is that of providing spatial Mercalli intensity reports based only Twitter data, we consider it beyond the scope of our current work to compare with approaches that use seismograph measurements. Regardless, we do believe it is important to understand the relationship between seismograph recordings and how Twitter data, in our case, can complement this data. We present further discussion on this in Sect. 5.

Table 10 shows the results of our proposal in relation to the work by Cresci *et al.* for the task of maximum intensity estimation. As mentioned earlier in Sect. 2, our method uses less features than the one by Cresci *et al.*. In addition, ours considers support and spatial smoothing.

Table 10 Testing averaged MAE of the maximum Mercalli intensity for each earthquake

	Proposal	Cresci <i>et al.</i> [9]	Baseline
2	(1.00,+1.0,-1.0)	(2.00,+0.0,-0.0)	(4.12,+0.6,-0.4)
3	(0.69,+1.3,-0.6)	(1.05,+0.9,-1.0)	(3.04,+1.0,-0.8)
4	(1.25,+2.7,-1.2)	(0.20,+0.7,-0.2)	(1.27,+1.4,-0.8)
5	(0.55,+0.4,-0.5)	(0.73,+1.2,-0.7)	(0.94,+0.8,-0.6)
6	(1.00,+1.0,-1.0)	(0.50,+0.5,-0.0)	(1.00,+1.0,-1.0)
7	(1.00,+0.0,-0.0)	(7.00,+0.0,-0.0)	(1.00,+0.0,-0.0)

As Table 10 reveals, our method performs well in the specific task of maximum intensity prediction, being competitive with the state of the art. The method of Cresci *et al.* [9] outperforms our method at intensity 4 but at the cost of much less accurate predictions for low- and high-energy seismic movements. Note that at intensity 7, the method of Cresci *et al.* cannot detect the earthquake, while our method reaches only 1 point in MAE. This noteworthy result is because this earthquake was located at a rural locality in Chile (Limache), an area of the country that is sparsely populated. This earthquake produced a local trend in the region of interest but was practically uncommented on in the capital of Chile during the first half hour before the event. Then, the event just produced a local trend in the region of interest that was successfully detected by our method and discarded by our competitor.

The improvement of our method over the baseline is important for low- and medium-energy events. Note that the baseline corresponds to a regression over the 12 lexical features at municipality level, picking the maximum value detected in each earthquake. Our proposal applies spatial smoothing over highly supported regressors to finally pick the maximum. The results show that spatial smoothing is also useful for the maximum intensity detection task of low- and medium-energy events.

4.3 Illustrative examples

In Table 11, we show some examples of Mercalli reports provided by the National Seismological Center in Chile^h and reports generated by our method. We show one earthquake (chosen at random) per level of intensity. To build each report, we included in the list the municipalities with the largest populations, excluding municipalities with fewer than 25,000 inhabitants. A similar procedure is used by the Seismological Center to construct its national reports. Each list was sorted in decreasing order according to Mercalli intensity.

Table 11 shows that our method works well, providing accurate spatial reports. The estimation of the maximum intensity is accurate, and it also detects the epicenter. Thus, our method can detect the maximum intensity of the earthquake as well as the location in which it was registered. The use of spatial smoothing gives excellent results in terms of damage detection along the national territory. As expected, high-intensity earthquakes produce longer reports, showing an almost perfect match with the actual report. The use of reinforced Mercalli support helps to detect medium- and low-energy events. Note that the earthquake at intensity 3, located in Talca, was successfully detected by our method and also matched for localities close to where the event was perceived.

5 Discussion and conclusions

In this paper, we propose the first method for predicting the distribution of spatial Mercalli intensities for earthquakes using only social media features. Our literature review shows

Table 11 Illustrative examples of spatial Mercalli reports comparing actual and predicted Mercalli intensities

M	Actual report	Predicted report
3	Talca (3), Constitución (3), Maule (3), Navidad (2), Santiago (1)	Maule (4), Talca (3), Constitución (3), Curepto (3)
4	Coquimbo (4), Ovalle (4), Melipilla (2), Santiago (2)	Coquimbo (3), Colina (3), Til-Til (3), Santiago (3)
5	La Serena (5), Coquimbo (5), Vicuña (4), Ovalle (4), Illapel (3)	La Serena (5), Coquimbo (5), Vicuña (4), Ovalle (4), Illapel (3)
6	Quintero (6), Valparaíso (5), Quilpue (5), Quillota (5), Viña del Mar (4), Ovalle(3), Santiago(3)	Quintero (6), Valparaíso (6), Viña del Mar (6), Quilpue (5), Quillota (5), San Felipe (4), Santiago (4), La Serena (3)
7	Limache(7), Santiago (6), Viña del Mar (6), Valparaíso (6), Coquimbo (5), La Serena (5), Ovalle (5), Rancagua (5), Curicó (4), Coronel (3)	Limache (6), Viña del Mar (6), Valparaíso (6), Santiago (5), Coquimbo (5), La Serena (5), Rancagua (5), Curicó (5), Ovalle (4), Quirihue (3)

many efforts towards earthquake detection using social media: mostly, the location where an earthquake was felt and the maximum earthquake intensity. Our proposal performs well for both the aforementioned tasks and is competitive in relation to the state-of-the-art. However, these are not the main goals of our work. Our main objective is to predict the spatial distribution of Mercalli intensities, without depending on geological models or on using signals captured by spatially distributed seismographs [36]. Our empirical evaluation shows that social media provides valuable spatial information, which is helpful for the task of producing spatial intensity reports for earthquakes. In addition, we were successful in revealing local trends by using local-level high-support regressors. Our method uses a fine level of granularity in its spatial analysis (as opposed to prior approaches that use a more coarse-grained analysis) which allows us to detect and provide reports for medium-energy and high-energy events. Our experimental results show that our estimated reports are almost identical to those produced by experts.

On the other hand, our approach is not without its limitations. The main restriction on our proposed method is its dependency on the availability of spatially distributed social media data, specifically from Twitter. We think it is possible to generalize our method use data from other social media platforms that contain textual messages and location information (given that this data is available). However, our method cannot work without some type of social media content. Therefore, in geographic places where there is little or no social media coverage we will not have enough data to produce accurate estimates. A similar situation can also occur when faced with disasters in which digital communications are interrupted and people are not able to post on social media. This type of limitations are not exclusive to our system, but they correspond to a drawback of all crisis informatics systems that rely solely on social media as a data source. Nevertheless, we do not see social media dependency as a threat, given that our system is designed for the purpose of providing social media information during crisis situations to enhance emergency response when possible.

In addition, another limitation is the quality of message location estimation. Currently, our approach uses a more or less standard heuristic method to infer message geolocation.

This is needed due to the lack of GPS data associated to Twitter usage in Chile (and in many other countries). This can induce noise in our location data, thus, by incorporating better geo-mapping techniques our method could improve its accuracy. Recent work shows promising results in this direction [38], proving that methods based on the user friendship network obtain high values of accuracy in geolocation inference. However, finding the best approach for location estimation is beyond the scope of our current work. On this note, we believe that our method can be scaled to other countries. This notion of scalability is supported by the fact that the existing literature provides more than sufficient evidence of the usefulness of Twitter in different countries (e.g., Italy, New Zealand, Australia, the U.S., the U.K., and Japan) for rapidly gaining situational awareness. Furthermore, countries such as the U.K. show a much wider adoption in the use of GPS enabled devices, which could actually improve the performance of our approach for these countries.

In summary, in this article we have presented a method for spatial inference of damages after an earthquake releasing results in the Mercalli scale. Our contribution is that of providing a tool to allow for early response and to improve coverage in locations where there are no expert observers. Gaining accurate situational awareness as soon as possible after a disaster is extremely valuable for emergency response agencies and governments.

Future work includes dealing with open issues, such as measuring and understanding the contribution of social media data in relation to other data sources such as seismograph recordings. We note that we do not expect Twitter to outperform methods that include seismograph measurements, which can be extremely accurate, but rather to study how the combination with social media can enhance report immediacy and quality. To achieve this one possibility would be combine features from both sources, as was done by Burks et al. [36]. Another open problem is that of studying our method's sensitivity to location estimation quality and also to the keyword-based approach used to retrieve relevant tweets. It is possible that these terms could induce message undersampling, since we could miss relevant tweets that do not include the selected terms. However, by including more keywords we also risk adding more noise to our dataset. In this sense, we the keywords that we are currently using are the same as in those used in [39], which have showed excellent recall for earthquakes in Chile. This is supported by Table 2, which shows that the selected terms allow us to achieve good coverage in our dataset for all relevant events in the country. Hence, we think that these terms offer a good trade-off between noise and relevant messages, given the linguistic variations found in Chile. Nevertheless, in other countries, where language is more diverse depending on the region, this could be an important limitation.

Additionally, in the future we contemplate extending our work to incorporate more features. Time-based features extracted from the Twitter stream (e.g., tweet interval rate) are a valuable source of information in the earthquake detection task. We believe these types of features can also be helpful in the elaboration of spatial intensity reports. In addition, since the reports produced by our method are almost identical to those produced by experts, we plan to embed *Twitter-based intensity estimations* into the state-of-the-art earthquake detection and visualization system Twicalliⁱ [39].

Acknowledgements

We thank Jazmine Maldonado from Inria Chile for her valuable help with our dataset curation. We also thank Hernan Sarmiento from the Universidad de Chile who helped us improve our coverage of the existing literature.

Funding

MM and BP acknowledge funding support from the Millennium Institute for Foundational Research on Data. MM was partially funded by the project BASAL FB0821. The funder played no role in the design of this study.

Abbreviations

GPS: global positioning system; GUC: geophysics University of Chile; MAE: mean absolute error; SMO: sequential minimal optimization; SVC: support vector classification; SVM: support vector machine.

Availability of data and materials

Data and its description are available at <https://doi.org/10.6084/m9.figshare.c.4206689>.

Competing interests

The authors declare that no conflicts of interest exist.

Authors' contributions

All authors collaboratively designed and performed the research, contributed new analytic tools, analyzed data, and wrote the paper. All authors read and approved the final manuscript.

Author details

¹Department of Informatics, Universidad Técnica Federico Santa María, Santiago, Chile. ²Department of Computer Science, Universidad de Chile, Santiago, Chile. ³Millennium Institute for Foundational Research on Data (IMFD), Santiago, Chile.

Endnotes

- ^a "Did you feel it?" website located at <https://earthquake.usgs.gov/data/dyfi/>.
- ^b <https://investor.twitterinc.com/>
- ^c <https://mstrohm.wordpress.com/2010/01/15/measuring-earthquakes-on-twitter-the-twicalli-scale/>
- ^d "Did You Feel It?" website of the U.S. Geological Survey (USGS).
- ^e *Fuzzy wuzzy*, a Python string matching library that uses the Levenshtein Distance to compare string sequences <https://github.com/seatgeek/fuzzywuzzy> (set to 80% fuzzy confidence level).
- ^f <https://developer.twitter.com/en/docs/tweets/search/api-reference/get-search-tweets.html>
- ^g <https://www.cs.waikato.ac.nz/~ml/weka/index.html>
- ^h <http://www.sismologia.cl>
- ⁱ <http://twicalli.cl/>

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 22 August 2018 Accepted: 24 January 2019 Published online: 31 January 2019

References

1. Castillo C (2016) Big crisis data. Cambridge University Press, Cambridge
2. Fajardo J, Yasumoto K, Shibata N, Sun W, Ito M (2014) Disaster information collection with opportunistic communication and message aggregation. *J Inf Process* 22(2):106–117
3. Lee Hughes A, Palen L (2009) Twitter adoption and use in mass convergence and emergency events. *Int J Emerg Manag* 6(3/4):248–260
4. Cameron MA, Power R, Robinson B, Yin J (2012) Emergency situation awareness from Twitter for crisis management. In: Proceedings of the 21st international conference on world wide web. WWW '12 companion. ACM, New York, pp 695–698
5. Imran M, Castillo C, Diaz F, Vieweg S (2015) Processing social media messages in mass emergency: a survey. *ACM Comput Surv* 47(4):1–38
6. Sakaki T, Okazaki M, Matsuo Y (2013) Tweet analysis for real-time event detection and earthquake reporting system development. *IEEE Trans Knowl Data Eng* 25(4):919–931
7. Avvenuti M, Cresci S, Marchetti A, Meletti C, Tesconi M (2014) EARS (earthquake alert and report system): a real time decision support system for earthquake crisis management. In: Proceedings of the 20th ACM SIGKDD international conference on knowledge discovery and data mining. ACM, New York, pp 1749–1758
8. Maldonado J, Guzman J, Poblete B (2017) A lightweight and real-time worldwide earthquake detection and monitoring system based on citizen sensors. In: Fifth AAAI conference on human computation and crowdsourcing (HCOMP 2017). AAAI Press, Menlo Park, pp 137–146
9. Cresci S, La Polla M, Marchetti A, Meletti C, Tesconi M (2014) Towards a Timely Prediction of Earthquake Intensity with Social Media. IIT TR-12/2014 Technical report, IIT: Istituto di Informatica e Telematica, CNR
10. Mendoza M, Poblete B, Valderrama I (2018) Early tracking of people's reaction in Twitter for fast reporting of damages in the Mercalli scale. In: Meiselwitz G (ed) *Social computing and social media. Technologies and analytics*. Springer, Berlin, pp 247–257
11. Zhou A, Qian W, Ma H (2012) Social media data analysis for revealing collective behaviors. In: Proceedings of the 18th ACM SIGKDD international conference on knowledge discovery and data mining. KDD '12. ACM, New York, pp 1402–1402
12. Earle P, Guy M, Buckmaster R, Ostrum C, Horvath S, Vaughan A (2010) OMG earthquake! Can Twitter improve earthquake response? *Seismol Res Lett* 81(2):246–251

13. Palen L, Anderson KM (2016) Crisis informatics—new data for extraordinary times. *Science* 353(6296):224–225. <https://doi.org/10.1126/science.aag2579>
14. Bagrow J, Wang D, Barabasi A (2011) Collective response of human populations to large-scale emergencies *PLoS ONE* 6(3):e17680
15. Mendoza M, Poblete B, Castillo C (2010) Twitter under crisis: can we trust what we RT? In: Proceedings of the first workshop on social media analytics. SOMA '10. ACM, New York, pp 71–79
16. Castillo C, Mendoza M, Poblete B (2013) Predicting information credibility in time-sensitive social media. *Internet Res* 23(5):560–588
17. De Longueville B, Smith RS, Luraschi G (2009) “OMG, from here, I can see the flames!”: a use case of mining location based social networks to acquire spatio-temporal data on forest fires. In: Proceedings of the 2009 international workshop on location based social networks. LBSN '09. ACM, New York, pp 73–80
18. Bauman K, Tuzhilin A, Zaczynski R (2017) Using social sensors for detecting emergency events: a case of power outages in the electrical utility industry. *ACM Trans Manag Inf Syst* 8(2–3):1–20
19. Steinert-Threlkeld Z, Mocanu D, Vespignani A, Fowler J (2015) Online social networks and offline protest. *EPJ Data Sci* 4(19)
20. Mukherjee T, Chander D, Eswaran S, Singh M, Varma P, Chugh A, Dasgupta K (2015) Janayuja: a people-centric platform to generate reliable and actionable insights for civic agencies. In: Proceedings of the 2015 annual symposium on computing for development. DEV '15. ACM, New York, pp 137–145
21. Vieweg S, Hughes AL, Starbird K, Palen L (2010) Microblogging during two natural hazards events: what Twitter may contribute to situational awareness. In: Proceedings of the SIGCHI conference on human factors in computing systems. ACM, New York, pp 1079–1088
22. Kryvasheyev Y, Chen H, Obradovich N, Moro E, Van Hentenryck P, Fowler J, Cebrian M (2016) Rapid assessment of disaster damage using social media activity. *Sci Adv* 2(3) <https://doi.org/10.1126/sciadv.1500779>
23. Li T, Zhou W, Zeng C, Wang Q, Zhou Q, Wang D, Xu J, Huang Y, Wang W, Zhang M, Luis S, Chen S-C, Risse N (2016) DI-DAP: an efficient disaster information delivery and analysis platform in disaster management. In: Proceedings of the 25th ACM international on conference on information and knowledge management. CIKM '16. ACM, New York, pp 1593–1602
24. Rehman FU, Afyouni I, Lbath A, Basalamah S (2017) Understanding the spatio-temporal scope of multi-scale social events. In: Proceedings of the 1st ACM SIGSPATIAL workshop on analytics for local events and news. LENS'17. ACM, New York, pp 1–7
25. Rosas E, Hidalgo N, Gil-Costa V, Bonacic C, Marin M, Senger H, Arantes L, Marcondes C, Marin O (2016) Survey on simulation for mobile ad-hoc communication for disaster scenarios. *J Comput Sci Technol* 31(2):326–349
26. Yin J, Lampert A, Cameron M, Robinson B, Power R (2012) Using social media to enhance emergency situation awareness. *IEEE Intell Syst* 27(6):52–59
27. Doan S, Vo B-K, Collier N (2011) An analysis of Twitter messages in the 2011 Tohoku earthquake. In: International conference on electronic healthcare. Springer, Berlin, pp 58–66
28. Murakami A, Nasukawa T (2012) Tweeting about the Tsunami?: mining Twitter for information on the Tohoku earthquake and Tsunami. In: Proceedings of the 21st international conference on world wide web. WWW '12 companion. ACM, New York, pp 709–710
29. Kropivnitskaya Y, Tiampo KF, Qin J, Bauer MA (2017) The predictive relationship between earthquake intensity and tweets rate for real-time ground-motion estimation. *Seismol Res Lett* 88(3):840–850
30. Crooks A, Croitoru A, Stefanidis A, Radzikowski J (2013) # Earthquake: Twitter as a distributed sensor system. *Trans GIS* 17(1):124–147
31. Robinson B, Power R, Cameron M (2013) A sensitive Twitter earthquake detector. In: Proceedings of the 22nd international conference on world wide web. ACM, New York, pp 999–1002
32. Zhang X, Shasha D (2006) Better burst detection. In: Data engineering, 2006. ICDE'06. Proceedings of the 22nd international conference on. IEEE Press, New York, pp 146–146
33. Sakaki T, Okazaki M, 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, New York, pp 851–860
34. D'Auria L, Convertito V (2016) Real-time mapping of earthquake perception areas in the Italian region from Twitter streams analysis. In: Earthquakes and their impact on society. Springer, Berlin, pp 619–630
35. Unankard S, Li X, Sharaf MA (2015) Emerging event detection in social networks with location sensitivity. *World Wide Web* 18(5):1393–1417 <https://doi.org/10.1007/s11280-014-0291-3>
36. Burks L, Miller M, Zadeh R (2014) Rapid estimate of ground shaking intensity by combining simple earthquake characteristics with tweets. In: 10th US nat. conf. earthquake eng., front earthquake eng. Anchorage
37. Yin J, Karimi S, Lingad J (2014) Pinpointing locational focus in microblogs. In: Proceedings of the 2014 Australasian document computing symposium. ACM, New York, p 66
38. Ribeiro S, Pappa GL (2018) Strategies for combining Twitter users geo-location methods. *Geoinformatica* 22(3):563–587
39. Poblete B, Guzmán J, Maldonado J, Tobar F (2018) Robust detection of extreme events using Twitter: worldwide earthquake monitoring. *IEEE Trans Multimed* 20(10):2551–2561. <https://doi.org/10.1109/TMM.2018.2855107>