Portability of semantic and spatial–temporal machine learning methods to analyse social media for near-real-time disaster monitoring

Clemens Havas¹ · Bernd Resch^{1,2}

Received: 21 March 2020 / Accepted: 19 May 2021 / Published online: 10 July 2021
© The Author(s) 2021, corrected publication 2021

Abstract

Up-to-date information about an emergency is crucial for effective disaster management. However, severe restrictions impede the creation of spatiotemporal information by current remote sensing-based monitoring systems, especially at the beginning of a disaster. Multiple publications have shown promising results in complementing monitoring systems through spatiotemporal information extracted from social media data. However, various monitoring system criteria, such as near-real-time capabilities or applicability for different disaster types and use cases, have not yet been addressed. This paper presents an improved version of a recently proposed methodology to identify disaster-impacted areas (hot spots and cold spots) by combining semantic and geospatial machine learning methods. The process of identifying impacted areas is automated using semi-supervised topic models for various kinds of natural disasters. We validated the portability of our approach through experiments with multiple natural disasters and disaster types with differing characteristics, whereby one use case served to prove the near-real-time capability of our approach. We demonstrated the validity of the produced information by comparing the results with official authority datasets provided by the United States Geological Survey and the National Hurricane Centre. The validation shows that our approach produces reliable results that match the official authority datasets. Furthermore, the analysis result values are shown and compared to the outputs of the remote sensing-based Copernicus Emergency Management Service. The information derived from different sources can thus be considered to reliably detect disaster-impacted areas that were not detected by the Copernicus Emergency Management Service, particularly in densely populated cities.

Keywords Social media · Disaster management · Machine learning · Semantic topic analysis · Geospatial analysis

✉ Clemens Havas
clemensrudolf.havas@sbg.ac.at

¹ Department of Geoinformatics, Paris-Lodron University of Salzburg, 5020 Salzburg, Austria

² Center for Geographic Analysis, Harvard University, Cambridge, MA 02138, USA

1 Introduction

In the last two decades, natural disasters have become more frequent and more severe, resulting in higher economic losses and death tolls (Wallemacq and House 2018). Due to climate change, this trend will continue, and information products from monitoring programmes like the European Union Copernicus Programme are essential to mitigate and prevent damages and losses caused by natural disasters. The Emergency Management Service (EMS) of the Copernicus Programme is deployed to provide timely and accurate geospatial information, mainly based on remote sensing satellite data. Although satellite imagery has proven to be a highly valuable data source for emergency management, remote sensing-based monitoring systems face certain limitations. For example, due to activation or orbital constraints, the temporal lag before the first information product provided by such a system can be up to 72 h after the outbreak of a disaster (Schnebele and Cervone 2013). Furthermore, the satellite's schedule and path can differ from the area of interest and the targeted period, making spatial and temporal data about a specific natural disaster sparse. Other limitations, such as orbital or physical constraints, can reduce the quality of the derived images.

A new EMS component has been designed to tackle these challenges by exploiting social media analysis and crowdsourcing capabilities to generate timely and accurate information (Havas et al. 2017). Messages shared on social media platforms such as Twitter data (Tweets) can complement remote sensing-based monitoring systems by providing information shared during the initial hours between when a disaster occurs and the remote sensing data become available. A significant advantage of social media data over traditional data is the near-real-time aspect of the data source, which is available without substantial temporal delays. Furthermore, social media data can be considered in situ sensor data because a user can provide location-specific insights about the local situation by describing things like damaged buildings or traffic blockades (Nguyen et al. 2017). Besides infrastructural problems, social media data is user-generated, leveraging specific local and contextual knowledge (Resch 2013). Humans share information and emotions on social media platforms that can provide valuable insights to disaster managers. Therefore, social media data does not only carry similar information as remote sensing data, but also represents how humans are impacted by a disaster and raises situational awareness in an emergency situation.

However, social media data (and user-generated data in general) is typically created without any mandatory form and includes a high degree of noise (Gao et al. 2011). Therefore, relevant data must be identified to extract useful information about an event, whereby the identification methods depend on the particular use case. For example, relevant data must be created within the period of the disaster and georeferenced in the area of interest to extract information about the temporal evolution of a natural disaster in an area. Still, the spatial and temporal dataset consists of unrelated data that require sophisticated methods to identify only the data relevant to the disaster. Fortunately, social media data is rich in features of multiple modalities that provide numerous information extraction options. Recent research efforts have contributed to social media analysis in the context of disaster management by analysing various content types such as text, images, videos, geolocation, time and user networks (Zhao et al. 2013).

An example where various analysis methods on different modalities of social media data are combined is the newly designed EMS component that offers end-to-end services, such as hot spot maps, based on the promising multimodal methodology presented by Resch et al. (2018). The methodology combines two machine learning methods to analyse textual

and geospatial features of social media data. Because text must be included in a social media post and the corresponding GPS coordinates are the most precise spatial attributes, those features are used as input for the machine learning algorithms. To extract disaster-related information from the social media dataset, two machine learning algorithms extract semantic topics related to a natural disaster from a spatially and temporally filtered dataset, and a geospatial hot spot analysis highlights high/low clusters of disaster-related social media posts. Latent Dirichlet allocation (LDA) was used to extract semantic topics from the Tweet corpus as it is a highly flexible unsupervised machine learning algorithm (Blei et al. 2003). The overall goal of this methodology is to identify areas impacted by natural disasters within a defined area of interest.

Although Resch et al. (2018) presented compelling results for the Napa earthquake in 2014, further experiments have not been conducted. Furthermore, they stated that the methodology could create delineation layers within 1–3 h, which has not yet been tested but constitutes a requirement for a near-real-time monitoring system. Another critical missing development step is the automatic interpretation of the semantic clustering of the Tweets that are created with topic models. The methodology requires manual supervision for both machine learning methods, which is not feasible for a near-real-time monitoring system. In this publication, we aim to improve the methodology by automating the processing of semantic topics. Additionally, we prove the portability of the methodology by testing it on natural disasters that differ in space and time. We also examine its portability to another disaster type, namely hurricanes. Two of the use cases ("cold case") are conducted with historical data and one in a currently ongoing disaster ("warm case") to test the near-real-time capabilities by continuously creating hot spot maps in the first days of an occurring hurricane and comparing the results with official authority datasets. The last step is to compare the results with the EMS outputs and official authority datasets to demonstrate the value of social media analysis for monitoring systems.

2 Related work

Effective disaster management relies on timely and accurate information about the impact of a natural disaster in a region. To create timely information, a continuous stream of data georeferenced in the area of interest must be analysed without prior knowledge about the areas impacted by the disaster. Numerous publications have shown the potential of social media data analysis for disaster management by analysing different content of a post (Middleton et al. 2014; Yoo et al. 2016; Yin et al. 2012; Reuter et al. 2013; Schnebele et al. 2014). While valuable information can be extracted from social media data, they must first be filtered as they can contain posts that are unrelated to the natural disaster.

Herfort et al. (2013) defined keywords associated with the Elbe flood event in Germany and extracted Tweets where a keyword is part of the Tweet's text. Likewise, De Albuquerque et al. (2015) and Zou et al. (Zou 2019) used keyword filtering to extract disaster-related Tweets. These authors all analysed the temporal and spatial distribution of the disaster-related Tweets and were able to determine the impacted areas. Other approaches that use keyword-based filtering in the context of disaster management categorised Tweets into "disaster-related" and "not disaster-related" or detect damage and casualty reporting (Guan and Chen 2014; Zhao et al. 2012). However, approaches based on manually defined

keywords are high maintenance and static, thus requiring more sophisticated machine learning methods that learn relevant keywords autonomously.

A topic model is a machine learning technique that learns keywords associated with a disaster. Topic models have been applied to Tweets in the context of disaster management. Kireyev et al. (2009) adjusted the LDA algorithm to enhance results for short text messages like Tweets by using a term weighting function. The authors produced valuable results for single events such as an earthquake or a tsunami but did not show the impact of their results in the geospatial domain.

To increase the number of georeferenced Tweets, Jongman et al. (2015) filtered Tweets by keywords and extracted locations from the Tweet's texts. In their publication, they analysed satellite observations and social media data in parallel for early flood detection. However, they neglected to consider the georeference of a Tweet, which is usually more precise than the locations extracted from the text. Similarly, Francalanci et al. (2017) extracted locations from Tweets based on keywords in combination with named entity recognition in the context of emergency management. They link the obtained locations to images to associate a precise location with coordinates to each Tweet. Although their algorithm increased the number of georeferenced Tweets, they were only able to geolocate a few Tweets. In contrast, in this publication, we provide a strategy to collect as much georeferenced social media data as possible.

3 Natural disaster events and data collection

In this section, we present the "geocrawler", which enables the collection of georeferenced posts from social networks. Furthermore, the datasets and use cases for the experiments are explained, and the time series and spatial distribution of Tweets are also shown.

3.1 Geocrawler

Natural disasters can occur unexpectedly, and their force makes them one of the most serious threats worldwide. In particular, earthquakes occur without warning, and hurricane tracks cannot be predicted accurately (Cox et al. 2013). To enhance analysis results and guarantee timely analysis, spatially and temporally relevant data must be collected as quickly as possible. We developed geocrawler software to collect as much relevant social media data as possible within a reasonable time by requesting data from the application programming interfaces (API) of social media platforms. The programme can query data from multiple social media platforms such as Twitter, Flickr, YouTube or Foursquare. As the focus of the experiments is on Twitter data, we only describe how we retrieve Tweets in this publication. However, it works similarly for other platforms with specific adaptations to each social network's application programming interfaces (API).

Twitter provides two types of APIs to collect Tweets: REST and streaming (Developer 2020). The REST API offers various endpoints for using the functionalities of Twitter, such as the endpoint "search/tweets", to collect Tweets from the last seven days with certain limitations, such as a maximum of 450 requests per 15-min interval. These constraints make the collection process challenging and require an advanced strategy to cope with the fast-moving time window of the API to harvest all offered Tweets with a minimal number of requests. Contrarily, the streaming API provides a real-time data stream that can be filtered

with multiple parameters. Our developed software focuses on natively georeferenced Tweets within an area of interest and uses Twitter's REST and streaming API (Fig. 1).

Combining the REST and the streaming API makes crawling robust against interruptions or backend issues that would lead to missing data. For example, if data from the streaming API cannot be stored in time or an interruption occurs, the missing data can be retrieved via the REST API, which provides Tweets from the last seven days. We believe that, with this strategy, we can collect as much data as possible. It is important to note that this software is designed to also request data from other social media platforms such as YouTube, Flickr and Foursquare if the appropriate credentials for the particular network are specified.

In the context of disaster management, the geocrawler starts when a natural disaster occurs or begins and requests data for the area of interest without additional configuration. The geocrawler can also be used for monitoring, as the streaming API complements the REST API restrictions that do not allow querying all regions sufficiently with only one developer account.

The geocrawler was used to collect the data for the use cases presented in this paper. The use cases represent different disaster types that did not occur in the same year or in the same region.

3.2 Emergency management service information products

The EMS provides information for emergency response related to different types of disasters as well as prevention, preparedness, response and recovery activities. It consists of two main components: rapid mapping and early warning. The mapping component has worldwide coverage and provides maps based on satellite imagery. Authorised users can send a Service Request Form directly to the European Response Coordination Centre to trigger the mapping service and request reference, delineation or grading maps. For the targeted use cases, mapping coverage is limited to the requested areas by authorised users. For the needs of this study, different EMS layers were used for the use cases:

- *Areas of interest (AOIs)* An area defined by the user, which guides and limits the production to specific areas considered by the user to be impacted by the event.
- *Event layer* A layer with the flood traces or damage information observed in the impacted areas using satellite data.

3.3 Official authority datasets

Our social media analysis results are validated with official authority datasets that show the actual measured areas impacted by a natural disaster. Two authorities that provide open-source data for natural disasters are the United States Geological Survey (USGS) and the National Hurricane Center. USGS is a scientific institution and part of the United States Department of the Interior (DOI). It provides reliable scientific information to minimise loss of life and property due to natural disasters as well as for other purposes such as managing water, biological, energy and mineral resources. The National Hurricane Center is part of the National Centers for Environmental Prediction, focusing

on hazardous tropical weather that includes forecasting and analysing the path and the impact of hurricanes.

3.4 Use cases

For big data analysis, publicly available Tweets are the most suitable social media data source as the majority of Tweets are shared publicly. For all use cases, we collected natively georeferenced Tweets that are located within the respective area of interest.

3.4.1 Amatrice earthquake

On 24 August 2016, an earthquake with a magnitude of 6.2 hit central Italy with the epicentre close to Accumoli, Italy (Geological Survey 2020a). The earthquake caused the death of more than 290 people and led to severe damage in the city of Amatrice and neighbouring cities (Abbott and Schiermeier 2016). We collected 48,992 georeferenced Tweets between 17 August and 9 September 2016 that were located within the bounding box [11.7°W, 41.7°S, 14.7°E, 43.8°N] in the World Geodetic System 1984 (WGS 84). The earthquake hit central Italy at 03:36:32 on 24 August 2016, causing a peak in the time series of the number of Tweets per hour, whereby the peak matches the hour when the earthquake hit (Figure 2). Such a sudden increase in social media data has also been observed in other publications, e.g. in Earle et al. (2012) and (Resch et al. 2018).

Fig. 1 Collecting Tweets with geocrawler

Fig. 2 Time series (UTC) of Tweets for the Amatrice earthquake

Resch et al. (2018) analysed the 2014 Napa Valley earthquake. To evaluate portability to other earthquake events in another global region, we analysed the earthquake in Amatrice, Italy. This is especially interesting as the official language in Italy is Italian as opposed to English as in the USA. Therefore, the Amatrice earthquake can give us insights into whether the methodology works in another language, for other natural disasters, and on another continent.

We identified the used languages in the Tweet corpus with the python package Polyglot (Al-Rfou 2015) and found that approximately 57.1% were written in English and approximately 15.4% were written in Italian. The used language could not be automatically determined for 23.6% of the Tweets. We expected more Italian Tweets than English Tweets, but the mixed distribution of languages still poses a challenge.

Due to the comparably large quantity of data, data points would strongly overlap and coalesce on a point map, making it impossible for the reader to draw meaningful conclusions. Spatial binning is a visualisation technique where data points are aggregated in shapes such as triangles, rectangles or hexagons. Hexagons have proven to be the most suitable shape to use (Battersby et al. 2017). Therefore, we aggregated Tweets in hexagon bins to show the spatial distribution of the Twitter datasets. Figure 3 shows the spatial distribution of Tweets for the Amatrice earthquake in the given period. Although the Amatrice earthquake was one of the most severe earthquakes in Italy, nearby urban centres show more activity than the area impacted by the natural disaster. Compared to other rural

Fig. 3 Spatial distribution of Tweets in central Italy

Fig. 4 EMS damage detection for damaged buildings in Amatrice

areas, the area around Amatrice shows more activity than other regions. However, further analysis is needed as unrelated Tweets should not be included in the results used for disaster management.

For the Amatrice earthquake, the EMS produced damage maps covering only a small part of the impacted area. As shown in Fig. 4, the focus of the EMS layer is on the area around Amatrice, and the red squares represent damaged buildings.

USGS provides a geospatial dataset of the peak ground acceleration (PGA) of the Amatrice earthquake. The most substantially impacted areas are close to the city Amatrice. The PGA represents the intensity of the earthquake and is used for the comparison in this analysis because the footprint obtained by the hot spot analysis of social media is ideally similar to the PGA areas.

Fig. 5 Time series (UTC) of Tweets for Hurricane Harvey

Fig. 6 Spatial distribution of Tweets in Texas

Fig. 7 EMS flood outlines related to Hurricane Harvey

3.4.2 Hurricane Harvey

Hurricane Harvey made landfall on 25 August 2017, close to Houston, Texas, which is home to two million inhabitants. As Hurricane Harvey was slowing down, heavy rainfall occurred in the area and caused horrendous damage, estimated at \$125 billion and 68 deaths (Geological Survey 2020b). We collected 135,723 georeferenced Tweets between 25 August and 7 September 2018 that were located within the bounding box $[-98.8^{\circ}\text{W}, 27.3^{\circ}\text{S}, -90.4^{\circ}\text{E}, 31.1^{\circ}\text{N}]$ in the World Geodetic System 1984 (WGS 84). Contrary to earthquakes, hurricanes do not cause a sudden increase in Tweets because they are slow moving and can be monitored days before landfall (Figure 5). By comparing the number of Tweets per hour in the week in which Hurricane Harvey occurred (blue) with the week before (light orange), we observe that the overall number of Tweets increased, but there is no visible peak.

Fig. 8 Aggregated high water marks related to Hurricane Harvey

Fig. 9 Time series (UTC) of Tweets on the East Coast of the USA in case of the Hurricane Florence

The spatial binning of georeferenced Tweets in the area impacted by Hurricane Harvey shows high activity in urban regions, especially in Austin, San Antonio, Houston and Baton Rouge. While Austin, San Antonio and Baton Rouge were not affected by Hurricane Harvey, most of the city of Houston was flooded. This visualisation cannot be used for disaster managers to draw conclusions about which areas are impacted by a disaster, and therefore, the dataset must be further analysed (Fig. 6).

The EMS was activated for specific regions in Texas to create flood outlines. As public authorities must request information for particular areas, some regions were not examined by EMS (Figure 7).

In the case of Hurricane Harvey, USGS collected high water marks as indicative evidence for the extent of flood inundation caused by Hurricane Harvey. A team of

Fig. 10 Time series (UTC) of Tweets in North and South Carolina in case of the Hurricane Florence

Fig. 11 Spatial distribution of Tweets on the East Coast

Fig. 12 Spatial distribution of Tweets in North and South Carolina

hydrologists and hydrologic technicians flagged and surveyed these marks for the Hurricane Harvey event on behalf of the US Federal Emergency Management Administration, which coordinates the response to a disaster in the USA. Overall, they collected 2,123 georeferenced high water marks distributed along the coast of Texas and around Houston, USA. The high water marks are densely spaced and, for visual comparison, were aggregated in a polygon shown in Fig. 8.

3.4.3 Hurricane Florence

On 14 September 2018, Hurricane Florence made landfall in North Carolina as a Category 1 hurricane. The main destruction took place in North and South Carolina and other states, and then, Hurricane Florence weakened and became a tropical storm. Overall, it caused up to \$22 billion in damage, and 51 people died as a result of Hurricane Florence (Borter 2018). From 12 September until 19 September, we collected 414,303 georeferenced Tweets that were located within the bounding box $[-85.1^{\circ}\text{W}, 24.1^{\circ}\text{S}, -70.5^{\circ}\text{E}, 41.1^{\circ}\text{N}]$ in the World Geodetic System 1984 (WGS 84). Similar to Hurricane Harvey, there is no significant peak in the time series, as would be expected for this type of natural disaster. The time series of the two areas of interests are highly similar and only differ in magnitude (Figs. 9, 10).

Fig. 13 EMS flood outlines in North and South Carolina

In Fig. 11, the spatial Tweet distribution shows high activity in urban areas and especially the area considered to be part of "Boswash" from Boston to Washington, D.C. (D. C. 1975). When we focus on the area where Hurricane Florence made landfall in Fig. 12, we again observe that urban areas have higher activity. Still, numerous places along the coast also show high activity. Like in the case of the Amatrice earthquake, the data must be analysed to provide useful information to disaster managers.

EMS was activated for several locations along the coast of North Carolina where Hurricane Florence made landfall. Although Hurricane Florence caused flooding in many AOIs, EMS was only able to delineate a smaller fraction of impacted areas (Fig. 13).

In the case of Hurricane Florence, USGS collected 769 high water marks. The high water marks are densely clustered, and for comparison, they were aggregated in a polygon shown in Fig. 14.

4 Methodology: extracting disaster-related information from social media posts

Georeferenced social media posts include very useful information for emergency services. However, social media posts contain unrelated data, and the high variance in quality also makes it challenging to identify relevant data. Therefore, the following preprocessing and machine learning methods are needed to extract relevant information.

4.1 Preprocessing

In natural language processing, numerous preprocessing steps improve the subsequent analysis methods. Especially in social media analysis, datasets contain a high percentage of irrelevant and noisy posts (Verma, et al. 2011). Preprocessing cleans the dataset but also reduces the processable information in the corpus. Preprocessing routines are designed with caution as they strongly impact the quality of the results of the subsequent analysis. Thus, the order and the selected preprocessing steps are critical for the analysis. In our analysis, the goal is to reduce the Tweet's content to the relevant words that can be related to an event. Every Tweet is preprocessed by transforming words to lowercase and removing unnecessary characters or words such as URLs, numbers, short words, unique words or

Fig. 14 High water marks in North and South Carolina

Fig. 15 Preprocessing workflow

Fig. 16 Fifteen most frequent words for Hurricane Florence

stop words. Additionally, the words are stemmed so that words in different forms with the same meaning are combined (Figure 15).

Due to the extensive preprocessing, numerous Tweets in the filtered dataset contain only a few words or are even empty. In Fig. 16, the 15 most frequent words are shown for the georeferenced Tweets from within the local area of interest in the Hurricane Florence use case. Besides the words "Carolina" and "Atlanta", none of the words can be directly associated with the use case, which shows the importance of filtering the social media corpus with an intelligent algorithm again.

4.2 Semantic information extraction

Machine learning methods can filter the dataset to extract only data that is related to the targeted event. A promising machine learning method is the latent Dirichlet allocation (LDA), which groups a dataset into topics (Blei et al. 2003). LDA offers a high degree of flexibility as multiple hyperparameters can be defined, such as the number of topics z , α and β (Figure 17). Therefore, LDA can be adapted to the size of the corpus. This makes it especially useful in the case of geo-social media posts because the geospatial distribution of social media posts differs spatially as social media use varies (Poushter et al. 2018).

The identified topics in the corpus are distributions over words that are controlled with the hyperparameter β . If β is set very low, the topic is represented only by a few words. As our analysis aims to have only one disaster-related topic, we assign a low value to the β parameter. The hyperparameter α controls the distribution of the topic-document distribution and, similarly to β , the number of topics for a document is low if α has a low value. The assumption is that microblogging messages such as Tweets, which can only contain up to 280 characters, cover only one particular topic (Nugroho et al. 2015).

When LDA models topics, a human has to interpret them manually. The hyperparameter z is difficult to set, and typically, the model must be executed numerous times until crisp topics are extracted. This makes the evaluation of various hyperparameter settings a tedious task and slows down the modelling process significantly. A more effective way is using a semi-supervised version of LDA such as GuidedLDA (Jagarlamudi et al. 2012; Ferner et al. 2020). In the case of GuidedLDA, seed words associated with different topics can be defined. Those seed words are incorporated as lexical priors into the topic model. In our experiments, we defined disaster-related seed words that are associated with the first topic. Seed words can also be defined for other topics, but in our case, we wanted to identify only the disaster-related Tweets. The seed words guide LDA in such a way that the first topic is more likely to be represented by predefined seed words, but this is not guaranteed. Due to the nature of LDA, other topics can still be represented by the seed words predefined for the first topic. The six most probable words (stemmed) of two topics in the case of Hurricane Florence are shown in Table 1 as examples that represent the disaster-related topic for the use case and another derived topic that is related to job advertisements.

Following the derived document-topic distribution of GuidedLDA, every Tweet can contain multiple topics. The document-topic distribution also shows the likelihood of a Tweet containing a particular topic. With the parameter α , we predefine that a Tweet consists only of a few topics, and, as Tweets predominantly only have one topic per Tweet, we classify each Tweet by choosing the topic with the highest probability for a Tweet.

4.3 Hot spot analysis

The georeferenced dataset is semantically filtered to include only disaster-relevant Tweets. However, the filtered dataset usually still has a number of Tweets that cannot be presented on a map without further processing.

To effectively help emergency services, a hot spot analysis is applied to the dataset. The results are statistically significant cold and hot spots produced using the Getis-Ord

Fig. 17 Graphic model of latent Dirichlet allocation (based on Griffiths & Steyvers, 2004)

Gi* statistic, which indicates spatial clusters of low or high values. Contextualising it to the targeted use cases, hot spots represent areas where disaster-related social media posts are clustered and cold spots show areas where disaster-related posts are scarce. In contrast to other analysis techniques, hot spot analysis includes the spatial neighbourhood. It is, therefore, more robust to outliers such as an uncommonly high social media activity by one single user.

To apply hot spot analysis, the unfiltered georeferenced social media posts, as well as the disaster-related posts, are aggregated in a fishnet. The fishnet is overlain with the area of interest that is estimated based on recent news reports in a warm case. The fishnet cell size depends on the number of social media posts and the size of the area of interest and is calculated as follows:

$$l = \sqrt{2 \frac{A}{n}}$$

where l is the side length of a grid cell, A is the size of the area of interest and n is the number of points in the area of interest (Wong and Lee 2005).

Due to the heterogeneous distribution of Twitter data between densely and sparsely populated areas, the relationship between disaster-related Tweets and the unfiltered Tweets is calculated for every cell. This follows the assumption that the information value of a large number of disaster-related Tweets in a rural area is higher than in an urban area where the usual social media activity is always higher. The relationship between disaster-related Tweets and unfiltered Tweets is used as the input for the hot spot analysis.

4.4 Validation datasets

The final result is a hot spot layer showing the impacted areas, which must be validated to demonstrate the value of the methodology. Therefore, we assess the quality of the hot spot layer by visually comparing the disaster-related footprints with official authority datasets that describe the areas impacted by the natural disaster (s. Section 3.3). Additionally, we compare the hot spot layers with the outputs of Copernicus EMS.

Table 1 Example of top words (stemmed) from a disaster-related topic and another topic for the used dataset in the case of Hurricane Florence

Disaster-related topic	Other topic
Carolina	Latest
Hurricane	Open
North	Job
Hurricaneflor	Click
Florence	Work
Report	Appli

Fig. 18 Hot and cold spots related to Hurricane Harvey

5 Results and validation: delineated areas based on social media

The used methodology applies two machine learning methods sequentially to extract valuable information for disaster managers to gain insights into the degree to which different regions are impacted by the disaster.

5.1 Semantic topics

The semantic "topic" derived by the topic model is created with a semi-supervised algorithm that guarantees meaningful results with appropriate parameters and datasets. Topics can help disaster managers to identify keywords and hashtags that are associated with the event.

GuidedLDA requires a few hyperparameters such as the number of topics t , the distribution parameters α and β and the predefined seed words for topics that guide the topic model. The number of topics t is empirically defined through previous experiences for the experiments where multiple values were tested, and a higher number of topics are chosen for a greater corpus as this follows the assumption that more topics are present in a greater corpus. The predefined seed words that seed the GuidedLDA procedure in our experiments are similar to other publications that identified representative keywords for earthquakes and hurricanes for similar studies (Zou 2019; Avvenuti et al. 2014). The seed words should be chosen to capture relevant information, to exclude irrelevant information and to be generalisable for other events. The seed words for the hurricane use cases are "hurricane", "flood", "storm" and either "Harvey" or "Florence". In the case of the Amatrice earthquake, the seed words are "terremoto", "tremor", "shake", "earthquake" and "Amatrice". It must be noted that although we use seed words in this analysis, the nature of LDA allows for discovering other disaster-related Tweets that do not contain one of the seed words in their text.

Fig. 19 Hot and cold spots related to the Amatrice earthquake

5.2 Spatial hot spots

Figure 18 illustrates the results of our spatial analysis of hurricane-related Tweets for the Hurricane Harvey use case. Most of the hot spots showing disaster-impacted areas are located along the coast and especially in Houston and Corpus Christi, whereas cold spots are detected in San Antonio, Austin, Lafayette and Baton Rouge.

Figure 19 shows the hot spot map for the Amatrice earthquake. The overlaid hot spots clearly identify the area around Amatrice as the disaster-impacted area. The epicentre of the earthquake was close to Amatrice, which matches the hot spots. Along the coast, in Rome, and in the area north-east of Amatrice, we identified cold spots showing areas that were not impacted by the disaster.

5.2.1 Warm case Hurricane Florence

On 12 September 2018, we received a notification from EMS that Hurricane Florence is approaching the East Coast of the USA, so we started the geocrawler to collect geo-referenced Tweets. Through the entire period of Hurricane Florence, we created new hot spot maps. The chosen area of interest impacts the analysis results, and therefore, we analysed two different AOIs: the US East Coast and the area where Hurricane Florence was expected to make landfall, i.e. the area of North and South Carolina. On 14 September, the

Fig. 20 Hot and cold spots related to Hurricane Florence in North and South Carolina

day Hurricane Florence made landfall, we created the first hot spot map, which is shown in Fig. 20. We determined the impacted area along the coast where the hurricane made landfall (represented by red cells).

Figure 21 shows the hot and cold spots for the East Coast, based on the georeferenced Tweets from 15 to 17 September. The hot spots show heavily impacted areas in North and South Carolina and can be correlated with the hurricane track until that point in time. This enables disaster managers to observe where a hurricane is moving and where the natural disaster impacts people.

Fig. 21 Hot and cold spots related to Hurricane Florence on the East Coast

5.3 Validation

5.3.1 Amatrice earthquake

In Fig. 22, we show the USGS earthquake footprint (coloured in shades of orange) in different PGA intensities to correlate the footprint with the obtained hot spot map. The hot spots (red cells) strongly correlated with the USGS earthquake footprint as the highest PGA values are in the same areas as the hot spots. Outside the strongly impacted areas of the earthquake, all of the cells are either not significant or show significantly low activity. Therefore, this can help disaster managers detect where people are impacted the most and plan their logistics accordingly.

In comparison with the EMS outputs, the social media hot spot maps delineate the areas most impacted by the earthquake, whereas the results of EMS identified damaged buildings. Hence, the hot spot maps provide disaster managers that use EMS with a new perspective on the event that cannot be derived by satellite-based information maps.

Fig. 22 Comparing USGS PGA with hot and cold spots for the Amatrice earthquake

Fig. 23 Comparing EMS flood outlines with hot and cold spots for Hurricane Harvey

Fig. 24 Comparing aggregated high water marks with hot and cold spots for Hurricane Harvey

5.3.2 Hurricane Harvey

The hot spot maps are compared with the EMS outputs and the aggregated high water marks that serve as the ground-truth dataset. Figure 23 illustrates the derived hot spot maps for Hurricane Harvey underlain with EMS's flood outlines (blue polygons). The hot spot layer also delineated most of the flooded areas that were delineated by EMS. However, EMS shows more flooded areas east of Austin. Whereas the hot spots show high activity in the urban area of Houston, flood outlines by EMS are only identified around Houston. Houston suffered immensely as a result of the hurricane, which flooded the entire inner city. The high buildings in Houston could lead to deflections that make it challenging for EMS to derive the flood outlines in urban areas. Furthermore, the city Corpus Christi south of Houston was strongly impacted by Hurricane Florence, which can be seen in the hot spot map but was not identified by EMS. EMS was not activated for this area and did not create flood outlines.

For disaster managers, Twitter-based hot spot maps help determine impacted areas, especially in urban areas, to complement EMS results. They also provide the disaster manager with a complete overview of the area of interest, which leads to the identification of flooded areas such as Corpus Christi that EMS would not have identified.

In Fig. 24, the hot spot layer is compared with the aggregated high water marks that serve as the ground-truth dataset. The hot spot layer strongly overlaps with the high water marks. The main difference between the two layers is in the areas east of Austin and around Corpus Christi, where they differ slightly. In the area east of Austin, there are no significant cells, showing that the number of disaster-related Tweets is neither high nor low in this area.

5.3.3 Hurricane Florence

Figure 25 shows the hot spot layer underlain with the EMS outlines. We zoomed in to the area of interest as the flood outlines by EMS are limited. The identified flood areas are concentrated along the coast and to the east of Charlotte. EMS did not identify many areas that the hot spot layer indicated as being impacted by flooding. By adding the aggregated high

Fig. 25 Comparing EMS flood outlines with hot and cold spots for Hurricane Florence

water marks in Fig. 26, the hot spot layer has a high overlap with the ground-truth dataset. Although EMS was activated for Charleston, it did not identify any flood outlines, while the hot spot layer and the aggregated high water marks show that it was flooded.

Hurricane Florence was a very slow-moving storm that weakening as it approached the coast. As a result, Florence's wind field did not penetrate inland as far as average, or faster-than-average moving storms typically do. In Fig. 27, the best-track wind swath is compared with the hot spot layer for the East Coast. The best-track wind swath is created by accumulating the best-track wind radii. The best wind radii represent an estimate of the maximum extent of the 34-, 50- and 64-kt wind radii in each of the four quadrants of the storm at six-h intervals. The hot spot layer and the best-track wind swath overlap along the coast, whereas the hot spot layers also show areas where extensive rain fell due to Hurricane Florence. Disaster managers can analyse the disaster on different spatial scales using hot spot layers to identify the most impacted areas in a particular area.

Fig. 26 Comparing aggregated high water marks with hot and cold spots for Hurricane Florence

6 Discussion

6.1 Temporal evolution of a natural disaster

Natural disasters such as earthquakes, hurricanes and forest fires impact areas differently over time. Dense temporal and spatial data are necessary to identify which areas are currently most severely affected. Compared to satellite imagery alone, including social media data provides enough data to determine the most affected areas, as indicated by our Hurricane Florence use case. Using our approach presented in this paper (based on (Resch et al. 2018)), we can monitor the change of the impacted areas over time, even when the hurricane became a storm and moved inland. The social media analysis and especially Tweets are such an important contribution to natural disaster monitoring because the large data volume allows for separately analysing different time periods in short update intervals. This is valuable information for disaster managers as they can prepare and change their logistics planning according to changes in the hot spot maps in

Fig. 27 Comparing hurricane swath with hot and cold spots for Hurricane Florence

near real time. Furthermore, new georeferenced data are continuously created and collected via APIs to gradually enhance the analysis results. However, one challenge is to remove “old” data from the analysed corpus as they could distort the identification of areas currently impacted by the event. This leads to the question of how to define the time period for which data are analysed, i.e. which data are considered “old” data and which data are considered “current” data in the sense of a near-real-time analysis.

6.2 Enhancing existing monitoring systems

Current disaster monitoring systems such as the Copernicus Programme can greatly benefit from the integration of alternative information sources such as social media. The additional value for the monitoring procedure is manifold: firstly, the hot spot map reflects a human-centred disaster map as humans generate (most of the) data on social media platforms. Moreover, when humans describe their situation during a natural disaster, topic modelling

algorithms extract single social media posts that are related to the disaster that may be useful to disaster managers (e.g. in-situ reports indicating the severity of a disaster through textual descriptions or image and video documents). Then, relevant posts can be spatially analysed, yielding valuable information in the form of hot spot maps that indicate spatial accumulations of disaster impacts together with their temporal development. However, disaster managers must consider that the GPS coordinates of a Tweet can deviate from the described content as social media users can describe events at other locations or an event in the past (Ogie et al. 2019).

Another advantage is the combination of multiple data sources in subsequent geospatial analysis to enhance the information available during a disaster. For instance, Zhang et al. (Zhang 2019) developed a decision support framework that combines various datasets such as Tweets and census data for emergency management. Similarly, remote sensing-based data can be combined with social media data and other data sources, whereby adequate automated semantic fusion approaches are yet to be developed.

Besides the hot spot maps, individual social media posts that are related to the disaster can also be identified through topic modelling approaches (see above). These relevant posts can contain valuable information in various formats, including links to other sources (e.g. URLs to Instagram images) or social media-based damage assessment at almost no cost (Enenkel et al. 2018). Rosser et al. (2017) fused images extracted from social media with satellite and topography data to generate a probability map indicating flooded areas on a pixel level.

Various previous publications have proposed the combination of social media data with remote sensing data. Although the richness of information has been shown, social media analysis is only being integrated into monitoring systems in a few cases.

6.3 Dependence on social media networks

Social media platforms control access to their data, and they have recently made their APIs more restrictive for academic use. They are commercially run networks and can thus potentially change the APIs at short notice for reasons of stringent commercial interest and thereby decrease the reliability of further use without additional contracts (Juhász et al. 2016). An example is Instagram, which consecutively shut down the main features of their API in 2018 and will eventually be unusable for developers. Active APIs may also change their functionality. For instance, Twitter recently announced the prospective removal of the option to tag Tweets with precise locations (e.g. the smartphone's current GPS position). This development—if enacted—would have a dramatic impact on social media analysis for disaster management.

Besides strategic changes, platforms may update their infrastructure, forcing developers to update their programs to request data from social media APIs. Furthermore, Twitter changed its business model and introduced paid APIs on a pay-per-use basis. Although the data access is better compared to the public API (access to all data rather than to a random sample), it is most likely too expensive for academic and non-commercial purposes. Lastly, a social media platform may shut down for various reasons, such as financial constraints or changes in corporate strategies, causing an API service to stop working altogether. All in all, social media are a highly valuable resource, but multiple uncertainties are coupled with their analysis and should be kept in mind when using it as a primary source in a disaster management system.

7 Conclusion

This article proved the portability of an improved methodology for identifying disaster-impacted areas of various use cases that are spatially and temporally different and have different disaster types. First, a recently proposed methodology was improved by substituting an unsupervised topic model with a semi-supervised topic model. This makes the methodology usable in an automated workflow, which is crucial for disaster management. After an initial use case, which analysed an earthquake in the USA, we tested our methodology for an earthquake that occurred in a non-English speaking country. We demonstrated that although the corpus consisted of Tweets of various languages, disaster-impacted areas (hot spots and cold spots) could be identified correctly and accurately. Then, we tested our methodology on another disaster type, a hurricane. The analysis showed that social media analysis complements a remote sensing-based monitoring system as the social media-based information layers show flooded areas that the EMS did not detect. After verifying the portability for these historical disaster events (“cold cases”), we also validated the methodology in a current situation (“warm case”) during Hurricane Florence. We were able to show that disaster-impacted areas can be delineated in near real time by creating 6-hourly hot spot maps. The results were validated with official authority datasets, confirming their accuracy and value for disaster managers.

This publication highlights the benefits and challenges of complementing existing monitoring systems like the EMS with social media-based information layers. The use of social media data results in various considerable benefits, such as higher spatial and temporal resolutions of information about a disaster event compared to remote sensing data alone. Another benefit is that geo-social media analysis works particularly well in densely populated city centres. Social media posts can also be considered in situ human reports and thus exhibit substantially different information compared to conventional monitoring systems using physical sensors.

Although the value of social media analysis for disaster management has been shown, the confidence in the results and their interpretation is still improvable, for instance, by using more social media data and data from other sources. For the conducted experiments, we also collected georeferenced Flickr posts and YouTube videos. However, due to their lower spatial and temporal resolution and the typically limited textual content, the used methodology and especially the semantic analysis must be adapted. Next, we plan to investigate how to fuse the various datasets to enhance the analysis results and generalise the used methodology for different social media sources. Furthermore, the fusion of georeferenced datasets, such as the data from the National Hurricane Center or the GDELT platform with social media data, may further improve the confidence of the analysis results. Lastly, identifying fake news or Tweets that describe events from other locations would enhance the results.

Acknowledgements We would like to thank Wendy Guan from Harvard University’s Center for Geographic Analysis for her support in our study.

Funding Open access funding provided by Paris Lodron University of Salzburg. We would like to express our gratitude to the European Commission’s Horizon 2020 program for supporting the project “E2mC—Evolution of Emergency Copernicus services”, reference number 730082. Furthermore, this research was funded by the Austrian Science Fund (FWF) through the Doctoral College GIScience at the University of Salzburg (DK W 1237-N23).

Declarations

Conflict of interests The authors declare no conflict of interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abbott A, Schiermeier Q (2016) Italian scientists shocked by earthquake devastation. *Nature* 537(7618):15–16
- Al-Rfou R (2015) “Polyglot: a massive multilingual natural language processing pipeline,” State University of New York at Stony Brook
- 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*
- Battersby SE, Strebe D, Finn MP (2017) Shapes on a plane: Evaluating the impact of projection distortion on spatial binning. *Cartogr Geogr Inf Sci* 44(5):410–421
- Blei DM, Ng AY, Jordan MI (2003) Latent Dirichlet Allocation. *J Mach Learn Res* 3:993–1022
- Borter G (2018) “Hurricane Florence death toll rises to 51,” [Online]. Available: <https://www.reuters.com/article/us-storm-florence/hurricane-florence-death-toll-rises-to-51-idUSKCN1MC2JJ>
- Cox J, House D, Lindell M (2013) Visualizing uncertainty in predicted hurricane tracks. *Int J Uncertain Quantif* 3(2):143–156
- De Albuquerque JP, Herfort B, Brenning A, Zipf A (2015) A geographic approach for combining social media and authoritative data towards identifying useful information for disaster management. *Int J Geogr Inf Sci* 29(4):667–689
- Twitter Developer (2020) “Twitter Developer Docs.” [Online]. Available: <https://developer.twitter.com/en/docs>
- Earle PS, Bowden DC, Guy M (2012) Twitter earthquake detection: earthquake monitoring in a social world. *Ann Geophys* 54(6):708–715
- Enenkel M, Saenz SM, Dookie DS, Braman L, Obradovich N, Kryvasheyev Y (2018) “Social media data analysis and feedback for advanced disaster risk management,” *arXiv Prepr. arXiv1802.02631*
- Ferner C, Havas C, Birnbacher E, Wegenkittl S, Resch B (2020) Automated seeded latent Dirichlet allocation for social media based event detection and mapping. *Information* 11(8):376
- Francalanci C, Guglielmino P, Montalcini M, Scalia G, Pernici B (2017) “IMEXT: A method and system to extract geolocated images from Tweets - Analysis of a case study,” in *Proceedings of the IEEE RCIS'17*
- Gao H, Wang X, Barbier G, Liu H (2011) “Promoting coordination for disaster relief—From crowdsourcing to coordination,” in *Social Computing, Behavioral-Cultural Modeling and Prediction*, pp. 197–204
- Griffiths TL, Steyvers M (2004) Finding scientific topics. *Proc Natl Acad Sci* 101(Suppl. 1):5228–5235
- Guan X, Chen C (2014) Using social media data to understand and assess disasters. *Nat Hazards* 74(2):837–850
- Havas C *et al.* (2017) “E2mC: improving emergency management service practice through social media and crowdsourcing analysis in near real time,” *Sensors*, 17(12)
- Herfort B, de Albuquerque JP, Schelhorn SJ, Zipf A (2014) “Does the spatiotemporal distribution of tweets match the spatiotemporal distribution of flood phenomena? A study about the River Elbe Flood in June 2013,” in *Proceedings of the 11th international ISCRAM conference*, pp. 747–751
- Jagaramudi J, Daumé H, Udupa R (2012) “Incorporating lexical priors into topic models,” in *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*

- Jongman B, Wagemaker J, Romero B, de Perez E (2015) Early flood detection for rapid humanitarian response: harnessing near real-time satellite and Twitter signals. *ISPRS Int J Geo-Information* 4(4):2246–2266
- Juhász L, Rousell A, JokarArsanjani J (2016) Technical guidelines to extract and analyze VGI from different platforms. *Data* 1(3):15
- Kireyev K, Palen L, Andersson KM (2009) “Applications of topics models to analysis of disaster-related twitter data,” in *NIPS Workshop on Applications for Topic Models: Text and Beyond.*, vol. 1
- Verma *et al.* (2011) “Natural language processing to the rescue? extracting” situational awareness” tweets during mass emergency,” in *Proceedings of the 5th International AAAI Conference on Weblogs and Social Media (ICWSM 2011)*
- Middleton SE, Middleton L, Modafferi S (2014) Real-time crisis mapping of natural disasters using social media. *IEEE Intell Syst* 29(2):9–17
- Miller DC (1975) *Leadership and power in the Bos-Wash Megalopolis*. Wiley, Netherland
- Nguyen DT, Ofli F, Imran M, Mitra P (2017) “Damage assessment from social media imagery data during disasters,” in *Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017*, pp. 569–576
- Nugroho R, Molla-Aliod D, Yang J, Zhong Y, Paris C, Nepal S (2015) “Incorporating tweet relationships into topic derivation,” in *Conference of the Pacific Association for Computational Linguistics*
- Ogie RI, Clarke RJ, Forehead H, Perez P (2019) Crowdsourced social media data for disaster management: lessons from the PetaJakarta. org project. *Comput Environ Urban Syst* 73:108–117
- Poushter J, Bishop C, Chwe H (2018) “Social media use continues to rise in developing countries but plateaus across developed ones,” [Online]. Available: <https://www.pewresearch.org/global/2018/06/19/social-media-use-continues-to-rise-in-developing-countries-but-plateaus-across-developed-ones/>
- Resch B (2013) “People as sensors and collective sensing-contextual observations complementing geo-sensor network measurements”, in *Progress in location-based services*. Springer, Berlin, pp 391–406
- Resch B, Usländer F, Havas C (2018) Combining machine-learning topic models and spatiotemporal analysis of social media data for disaster footprint and damage assessment. *Cartogr Geogr Inf Sci* 45(4):362–376
- Reuter C, Heger O, Pipek V (2013) “Combining real and virtual volunteers through social media.,” in *Iscram*, pp. 780–790
- Rosser JF, Leibovici DG, Jackson MJ (2017) Rapid flood inundation mapping using social media, remote sensing and topographic data. *Nat Hazards* 87(1):103–120
- Schnebele E, Cervone G (2013) Improving remote sensing flood assessment using volunteered geographical data. *Nat Hazards Earth Syst Sci* 13(3):669–677
- Schnebele E, Cervone G, Waters N (2014) Road assessment after flood events using non-authoritative data. *Nat Hazards Earth Syst Sci* 14(4):1007–1015
- U.S. Geological Survey (2020a) “Magnitude 6.2 Earthquake in Central Italy.” [Online]. Available: <https://www.usgs.gov/news/magnitude-62-earthquake-central-italy>
- U.S. Geological Survey (2020b) “USGS response to hurricane Harvey—inundation maps and flood details.” [Online]. Available: <https://webapps.usgs.gov/harvey/>
- Wallemacq P, House R (2018) “Economic losses, poverty & disasters: 1998–2017,”
- Wong DWS, Lee J (2005) *Statistical analysis of geographic information with ArcView GIS and ArcGIS*. Wiley, Hoboken
- 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
- Yoo E, Rand W, Eftekhari M, Rabinovich E (2016) Evaluating information diffusion speed and its determinants in social media networks during humanitarian crises. *J Oper Manag* 45:123–133
- Zhang Z *et al* (2019) A cyberGIS-enabled multi-criteria spatial decision support system: a case study on flood emergency management. *Int J Digit Earth* 12(11):1364–1381
- Zhao J, Dong L, Wu J, Xu K (2012) “Moodlens: an emoticon-based sentiment analysis system for chinese tweets,” in *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1528–1531.
- Zhao X, Zhu F, Qian W, Zhou A (2013) “Impact of multimedia in sina weibo: popularity and life span”, in *Semantic Web and Web Science*. Springer, New York
- Zou L *et al* (2019) Social and geographical disparities in Twitter use during Hurricane Harvey. *Int J Digit Earth* 12(11):1300–1318