

# TriggerCit: Early Flood Alerting using Twitter and Geolocation - a comparison with alternative sources

**Carlo Bono and Barbara Pernici\***

Politecnico di Milano - DEIB  
[name.lastname@polimi.it](mailto:name.lastname@polimi.it)

**Jose Luis Fernandez-Marquez and Amudha Ravi Shankar**

University of Geneva

**Mehmet Oğuz Mülâyim**

Artificial Intelligence Research Institute  
(IIIA-CSIC)

**Edoardo Nemni**

United Nations Satellite Centre (UNOSAT),  
United Nations Institute for Training and  
Research (UNITAR)

## ABSTRACT

Rapid impact assessment in the immediate aftermath of a natural disaster is essential to provide adequate information to international organisations, local authorities, and first responders. Social media can support emergency response with evidence-based content posted by citizens and organisations during ongoing events. In the paper, we propose TriggerCit: an early flood alerting tool with a multilanguage approach focused on timeliness and geolocation. The paper focuses on assessing the reliability of the approach as a triggering system, comparing it with alternative sources for alerts, and evaluating the quality and amount of complementary information gathered. Geolocated visual evidence extracted from Twitter by TriggerCit was analysed in two case studies on floods in Thailand and Nepal in 2021.

## Keywords

Social Media, Disaster management, Early Alerting

## INTRODUCTION

The increasing use of mobile phones and social media has transformed the way people witness and report climate and weather related disasters. Social media has been demonstrated to be a worthy source of data to provide situational awareness in the course of a disaster event. Effective response to natural hazards, such as floods, can reduce loss of life and mitigate structural damages. Access to timely and accurate data is essential to the disaster response. The effectiveness of the humanitarian response is outside the scope of this work. Instead, we focus on the mechanism of early alert using social media, under the assumption that supporting the activation of the disaster response protocol with pertinent and geolocated evidence is beneficial for the affected community.

Remote sensing data such as satellite images have also proven to be a valuable source of information to support humanitarian relief and disaster response. Advances in Artificial Intelligence (AI) provide a way of extracting relevant information from satellite imagery and thus contribute to speeding up the damage assessment. Even though computer vision techniques can process large amounts of data in a short time, the triggering of such systems is often tied to manual or semi-automated procedures.

To leverage an alternative and abundant source - social media - this paper proposes a methodology for early flood alerting using Twitter, providing a systematic approach to extract near-real-time and geolocated evidence about

---

\*corresponding authors

onsets of flood events. The goal of this paper is to analyse how this information correlates with existing alert mechanisms.

The paper is structured as follows: the next section presents existing disaster alert and trigger mechanisms, and the current use of social media as an early alerting system. In the subsequent section, the TriggerCit methodology for extracting information from social media posts is illustrated. Then, the TriggerCit methodology is applied to two different case studies for flood events in Thailand and Nepal and the results are discussed.

## RELATED WORK

### Existing disaster alert and trigger mechanisms

Multi-hazard early warning systems and disaster alert mechanisms are a significant component of disaster risk reduction. They help prevent loss of life, reduce the impact of natural hazards, and assist responders in disaster handling. In this section, we will discuss two key state-of-the-art alert mechanisms.

*The Global Disaster Alert and Coordination System*<sup>1</sup> (GDACS), started in 2004, is a joint initiative of the United Nations Office for the Coordination of Humanitarian Affairs (UN OCHA)<sup>2</sup> and the European Commission (EC). It is designed to alert the international community during sudden-onset disasters that might require international assistance. GDACS serves to consolidate and improve the dissemination of disaster-related information and coordination in the first phase of a disaster. The GDACS website consists of different tools, which have been incorporated into disaster response plans by several governments and organisations. The GDACS Disaster Alerts are auto-generated based on algorithms and available data to alert oncoming disasters. The system calculates alerts based on specific risk indicators, like the severity of the disaster, the affected population, and the country's vulnerability. Subscribers are notified via Email and SMS (Kavallieros et al. 2015). GDACS information is openly accessible. For the first phase of a disaster, governments and disaster response organisations rely on GDACS alerts and automatic impact estimations to plan international assistance. The Virtual On-Site Operations Coordination Center (Virtual OSOCC)<sup>3</sup> facilitates information exchange and coordination.

*The Global Flood Awareness System*<sup>4</sup> (GloFAS) is a component of the Copernicus Emergency Management Service (CEMS) that delivers global hydrological forecasts to registered users. GloFAS combines information from the satellites and forecast models to produce GloFAS forecasts, Seasonal forecasts, and Impact Forecasts (Kavallieros et al. 2015; Zajac et al. 2014). The GloFAS Forecasts are available free of cost to all, with registered users from Hydro-Meteorological services, academia, and the humanitarian sector (*HydroConference | WMO Accessed on 01/21/2022*).

### Social media as early warning and alerting systems

Several social media sources have been considered in the literature, the main one being Twitter<sup>5</sup> as it provides public application programming interfaces (APIs) to search posts and related metadata. Several research works have been published on tweets analysis. Here we focus on tweet analysis for early event detection. One of the first papers on the topic is Sakaki et al. 2013, in which the focus is on the timeliness of the extracted information, in particular in the case of earthquakes. Like in other papers analysing social media, the post geolocation is extracted from Twitter metadata. As studied in the literature (e.g., Scalia et al. 2021), this is usually available only in a small percentage of the posts. Recent papers focused on analysing in parallel several sources of information for deriving awareness on emergency events. In Shoyama et al. 2021, the impact of extensive floods in Japan has been assessed using hydrogeological information and tweets for detecting event outbreaks and phases in the affected areas.

The importance of images and geolocation via text has been recently emphasised to overcome the limitations of native Twitter geolocations. There are several advantages of having precisely geolocated tweets. Fohringer et al. 2015 show how a very limited number of tweets, combined with elevation profiles, can give a precise map of flooded areas; also Scotti et al. 2020 emphasise the integration of tweet data and hydraulic modeling for mapping flooded areas. Wang et al. 2020 uses text to geolocate posts and to classify phases of events automatically. The geolocation steps are based on Named-Entity Recognition (NER) and gazetteers to identify locations. However, the focus is mainly on locations in the US and the English language.

---

<sup>1</sup><https://www.gdacs.org>

<sup>2</sup><https://www.unocha.org>

<sup>3</sup><https://vosocc.unocha.org>

<sup>4</sup><https://www.globalfloods.eu>

<sup>5</sup><https://twitter.com>

A more comprehensive approach to social media posts geolocation has been proposed in the CIME system (Scalia et al. 2021). The NER is based on Polyglot (Al-Rfou et al. 2015), which supports more than 40 languages. CIME disambiguates candidate locations using Nominatim<sup>6</sup>, a search engine for OpenStreetMap data. Disambiguation considers candidate locations, their distances, and their ranks in the OpenStreetMap administration level hierarchy.

Post geolocation can achieve a significant, although variable, level of precision. CIME was exploited in the E2mC project (Havas et al. 2017) to support rapid mapping activities within Copernicus Emergency Management Services (EMS)<sup>7</sup>. An important issue that emerged is the high number of duplicate and non-relevant images. A systematic approach to image filtering for geolocated tweets has been proposed in Negri et al. 2021, developing the VisualCit image processing toolkit, aimed at extracting indicators about emergency events. The toolkit allows image filtering based on contents, such as similarity, non-photographic content, presence of specific objects, and event-specific classifiers (including floods, see Rivera et al. 2021). Other authors have proposed to create event-specific classifiers (e.g., Imran et al. 2020) automatically, without focusing on locations and systematic data preprocessing.

Another important issue is determining the severity of events using tweets as a source of citizen knowledge. A combination of sentiment analysis and frequency of impact-related keywords has been used in Kankanamge et al. 2020.

The availability of a representative set of search keywords is paramount for leveraging social media data. The goal is to get sufficient coverage of the event, both in the initial phase and as the event unfolds, as shown in Olteanu et al. 2014. In several papers, general purpose keywords are used. A systematic approach has been developed in E2mC based on data mining techniques for multilingual keyword extraction (Kirsch et al. 2018).

In the present paper we propose TriggerCit, a systematic methodology to extract flood alerts from Twitter posts. We focus on detecting the occurrence of flood events within the first 48 hours of the event, and describing them at the regional administrative level, providing image evidence with detailed locations. The goal is to select localised, relevant information from social media in a short time. The approach leverages VisualCit processing and CIME geolocation. The results are compared, over two case studies, with data from the event alert systems described in 3.1.

## METHODOLOGY

The proposed approach, shown in Fig. 1, is implemented as an analytics pipeline aimed at detecting flood events and extracting high-quality event-related information. The components are adapted to Twitter. However, most of the concepts can be generalised to other compatible social media.




Figure 1. TriggerCit methodology for identifying occurrences of flood events and related locations.

The main steps of TriggerCit are illustrated in the following.

<sup>6</sup><https://nominatim.openstreetmap.org>

<sup>7</sup><https://emergency.copernicus.eu>

### 1. *Word dictionaries for social sensing*

A *seed dictionary*, which consists of a curated dictionary of words, is built for each language and event of interest. Ideally, this dictionary should be rich enough to account event variability, yet specific enough to match relevant messages only. Language-dependent dictionaries should be curated by expert native speakers. Moreover, the same event could be described in different languages, since multiple languages could be spoken in affected regions. In the current experiments, we assume that a small general-purpose set of terms is available, as we provide novel techniques for augmenting the event-related set of words (see Step 7).

### 2. *Word count monitoring for event detection*

With proper linguistic resources in place, a monitoring system for disaster-related terms can be implemented. Dictionaries can be used to build one or more queries towards this goal. We chose to build a query as the logical OR of the words in the dictionary. Alternatively, per-word time series can be used. Getting the time series of the number of matched posts is straightforward with Twitter APIs. The main advantage of querying counts is that no actual tweet content data is exchanged, so almost no computation is involved at this stage. Moreover, no tweet quota is consumed, which could be critical if the number of queries to monitor is high (e.g., many countries/languages). Rate limits are compatible with near-real-time monitoring. Geographical boundaries can be specified: however, this will limit the results to the few tweets that are natively geotagged in Twitter, which is usually a limiting factor as discussed in Scalia et al. 2021 and other works.

### 3. *Triggering*

Using the counts time series, an event triggering system can be designed. Ideally, the component should be a supervised classifier, using one or more validated datasets such as historical GDACS alerts for the event starts (and ends) as reference data. A more straightforward mechanism triggers an event if the tweet count growth exceeds some threshold.

### 4. *Get tweet contents*

If an event is likely occurring, we selectively download tweet data for the time interval of the event using the Twitter search endpoint, excluding retweets. Full-archive search API is available for research purposes. The recent search API is also available, which is accessible to the general public and limited to the last 7 days.

### 5. *Extract relevant content with VisualCit*

We then concentrate on tweets with images, applying a VisualCit semantic filtering pipeline to extract relevant contents. The tasks in the pipeline are ordered “quickest first” as a heuristic to get a faster overall processing time. The proposed pipeline consists of the following tasks:

- (a) Remove duplicated images and similar images (e.g., rescaled versions, modified versions).
- (b) Remove non-photos, such as drawings, screenshots and computer-generated images.
- (c) Remove not-safe-for-work images.

The confidence thresholds for the single tasks are chosen through empirical validation, to achieve a reasonable performance without substantially removing useful data.

### 6. *Use CIME for geolocation*

The CIME geolocation algorithm described in Scalia et al. 2021 is then applied to associate tweets with geographical entities, whenever possible. Locations in unrelated countries are filtered out using national administrative boundaries.

### 7. *Use pipeline output as feedback for input*

Having a selected set of relevant items could be useful to leverage the significant, unfiltered mass of tweets without images. A straightforward approach could be extracting tweets related to the ones in the VisualCit output, searching in the broader set of text-only tweets retrieved in Step 4. Textual frequency analysis may help us to identify keywords that were not explicitly listed in the original seed dictionary. New keywords can then be posed as a query to the corpus of text-only tweets. A simple measure like the TF-IDF score (possibly combined with the widespread cosine similarity) can rank individual tweets concerning this query. Selected text-only tweets with a score higher than some threshold (henceforth referred to as *promising text-only tweets*) can then be fed to CIME for geolocation, thus expanding the real-time geographical description of the event. In next section, focusing on the use case in Thailand, we show how this approach can effectively augment the map of relevant tweets. This approach could be further extended in order to refine the input space by recovering information with additional queries in Step 4.

### 8. Visualization

The process output consists of a representative set of social media posts, used to deliver thematic maps providing geolocated evidence to stakeholders.

## CASE STUDIES

In this section, we apply the proposed methodology to two flood events that occurred in 2021 in Thailand and Nepal. The two events were selected since the United Nations Satellite Centre (UNOSAT)<sup>8</sup> supported both activations in Nepal<sup>9</sup> and Thailand<sup>10</sup> with satellite-derived maps and AI-based flood detection analysis, updating operational dashboards in near-real-time<sup>11</sup> and providing daily updates.

### Thailand Flood Event - September-October 2021

#### Event description

In late September and October 2021, tropical storm Dianmu inundated regions in Vietnam, Laos, Cambodia, and Thailand. According to the UNOSAT Thailand flood monitoring dashboard<sup>12</sup>, around 14.3 square kilometers of the country were inundated between September 27 and October 27, 2021. An estimated 1.4 million people were affected by the flooding, with 32 of 76 provinces affected by heavy rains for nearly a month. The statistics mentioned above were computed in the aftermath of the flood events and should be considered as a preliminary analysis since they have not been validated yet in the field.

#### Data description

We queried Twitter with the logical OR of the three most frequent terms in the dictionary (see Fig. 2(a)) of Thai terms correlated to flood events and provided by the UNOSAT office in Bangkok. The event onset is clear using the Thai dictionary, but not with the equivalent English dictionary (see Fig. 2(b)).

We concentrated the analysis on the 48 hours between September 26, 00:00 UTC and September 28, 00:00 UTC, since there is marked evidence of the event onset. Over this period, more than 4 million tweets match the seed query.

Data	Count
All tweets, September 26-27	4'145'447
No retweets	66'868
Containing images	6'292
Native Twitter locations	227
Overall images	8'774
Passed VisualCit filters	3'056
Places geolocated by CIME	1'671

**Table 1. Dataset cardinality through processing, Thailand case study**

Level	Admin level	Count
Province	4	7
District	6	8
Municipality or subdistrict	8	176
Village or community	10	9
Other points	15	1'265

**Table 2. Administrative level counts for CIME geolocations**

#### Processing output

We first focused on processing the 8'774 images available<sup>13</sup> with VisualCit (see Fig. 4 for an example of the result).

<sup>8</sup><https://unitar.org/sustainable-development-goals/united-nations-satellite-centre-UNOSAT>


<sup>9</sup><https://unitar.org/maps/countries/70>

<sup>10</sup><https://unitar.org/maps/countries/100>


<sup>11</sup><https://unitar.org/about/news-stories/news/unosat-flood-ai-dashboards-nepal-creation-one-stop-shop-real-time-evidence-based-decision-making>

<sup>12</sup><https://unosat-geodrr.cern.ch/portal/apps/opsdashboard/index.html#/4f878691713a40f3b8ef3140e63c9f6d>

<sup>13</sup>The datasets for this paper will be made available on Zenodo.




(a) Thai entries



(b) Thai / English

**Figure 2. Daily tweet counts for Thai dictionary entries**

The VisualCit pipeline is configured to remove duplicates, non-photos, and not-safe-for-work contents. The average processing time for each item, including network overheads, is 171ms (roughly 20'000 elements/hour) on a single server.




(a) Native Twitter geolocation (b) VisualCit geolocation (c) VisualCit geolocation with extended tweet set (d) Affected people

**Figure 3. Geolocations / inhabitants ratio by region (a) Twitter native geolocations, (b) Twitter native + CIME geolocated, (c) Twitter native + CIME geolocated from extended dataset (with images + promising text-only tweets), and (d) Number of affected persons by region at September, 28th (source: ReliefWeb)**

The resulting dataset is then processed with the CIME algorithm in order to extract locations. The distribution of the places extracted by CIME does not simply replicate the distribution of the native geotags, as it can be seen comparing Fig. 3(a) and Fig. 3(b), and allows a clearer description of affected areas. In our experimental setup, to the 206 locations natively present in the tweets and located in Thailand, we can add 1'465 locations that are

within Thailand state borders. The average processing time for each item with CIME is 224ms<sup>14</sup>. The correlation between the geolocated data and the affected areas, as described by official sources<sup>15</sup>, can be appreciated in Fig. 3. Population estimates by province used for data normalisation were extracted from Wikipedia<sup>16</sup>. Mainly, specific locations or entities with sub-district administrative level are extracted (see Table 2<sup>17</sup>). Data cardinalities through processing steps are described in Tab. 1.

Fig. 3(c) shows how we can further extend location analysis by exploiting the text-only tweets as described in Step 7 of the TriggerCit methodology. A word frequency analysis of the VisualCit output provided us with two new keywords (water and house). We posed these two together with the original seeds as a query to the text-only corpus and measured the cumulative TF-IDF score of each tweet. We regarded the 6'021 tweets within the 0.9 percentile of the score as *promising* and fed them to CIME for geolocation. 661 of these tweets were successfully geolocated within Thailand by CIME.



**Figure 4. Example images for (a) items being removed by the pipeline, and (b) items being kept**

#### Assessment of the results

TriggerCit outputs an alert on September 27 due to the increase of tweets compared to the days before considered as a baseline. It also returns a weighted set of administrative levels where the alert is coming from, which partially correlates with the number of affected people and gives a preliminary indication of the affected regions.

#### Nepal Flood Event - June-July 2021

##### Event description

Between June and October 2021, 673 people have lost their lives in Nepal due to floods and landslides caused by seasonal and unseasonal heavy rain and severe weather<sup>18</sup>. Among the events at the beginning of the monsoon season, heavy rain triggered multiple landslides and flooding on the Melamchi river, in the Sindhupalchok district, resulting in more than 20 deaths<sup>19</sup>. Damage and casualties followed since late June and the beginning of July, marking the onset of the disaster event. The case study focuses on evidence from tweets gathered in this early stage.

##### Data description

Using a pre-defined dictionary of Nepali words generically related to flood events, the query consisted of the logical OR of the terms (see Fig. 5). As in the case of Thailand floods, it is difficult to extract evidence with English seed keywords (see Fig. 5(c)).

The most evident event using our dictionary happened in the middle of June, as seen in Fig. 5(a), which coincides with the event dates reported in the news. Over June 16-17, 6'639 tweets matched the query. It is relevant that, among the few posts with native Twitter geolocations, none contained information specific to this event.

We replicated the analysis for the first two days of July, since in these days several incidents with casualties happened in different places of Nepal<sup>20</sup>. In this case, the time series of tweets matching our query show an increasing trend, albeit less evident than in the previous case.

<sup>14</sup>With no parallelism

<sup>15</sup>[https://reliefweb.int/sites/reliefweb.int/files/resources/FlashUpdate\\_02\\_28Sep2021-TC-DIANMU-THLVNMLAO.pdf](https://reliefweb.int/sites/reliefweb.int/files/resources/FlashUpdate_02_28Sep2021-TC-DIANMU-THLVNMLAO.pdf)


<sup>16</sup>[https://en.wikipedia.org/wiki/Provinces\\_of\\_Thailand](https://en.wikipedia.org/wiki/Provinces_of_Thailand)

<sup>17</sup>OpenStreetMap administrative levels for Thailand <https://wiki.openstreetmap.org/wiki/Tag:boundary=administrative>.


<sup>18</sup><https://reliefweb.int/disaster/fl-2021-000134-npl#overview>

<sup>19</sup>[https://en.wikipedia.org/wiki/2021\\_Melamchi\\_flood](https://en.wikipedia.org/wiki/2021_Melamchi_flood)


<sup>20</sup><https://floodlist.com/asia/nepal-floods-landslides-july-2021>



(a) Nepali entries in June



(b) Nepali entries in July



(c) Nepali / English

**Figure 5. Daily tweet counts for Nepali dictionary entries**

### Processing output

The experimental setup is equivalent to the one used for the Thailand case study. Data cardinalities through processing are reported in Tab. 3. Since the pipeline did not support Nepali natively, we processed Nepali texts with a publicly available NER tool for Nepali<sup>21</sup>, described in Singh et al. 2019. The processed posts are then fed to a modified CIME algorithm for location extraction.

For data in June 16-17, among the 51 locations extracted, 21 refer to the Bagmati province, in which the event is taking place, and 19 explicitly refer to Sindhupalchok district (as Fig. 6 shows). Another 11 generically refer to


<sup>21</sup><https://github.com/oya163/nepali-ner>

Data	June 16-17	July 1-2
All tweets	6'639	1'225
No retweets	2'807	594
Containing images	261	63
Native Twitter locations	8	10
Overall images	391	80
Passed through VisualCit	218	55
Places geolocated with CIME	51	10

**Table 3. Data cardinality in Nepal case study for June 16-17 and July 1-2**

Nepal, so they are ignored. It is interesting to note that, compared to the GloFAS Rapid Impact Assessment map, there is a complementarity regarding this event, since there are no reporting points for the Sindhupalchok district.

For data in July 1-2, an overview of the processed data is reported in Table 3. The amount of evidence gathered is smaller than the previous time frame, with only 6 useful locations extracted after running the images through the pipeline (4 out of 10 are generic mentions of Nepal), and 14 in the overall set of tweets with images. Nevertheless, it is worth comparing the results with different sources of information to understand how to complement available information, which is not always concordant and could suffer from data scarcity. In Fig. 7 we gathered figures for different signals on July 2, coming from UNOSAT, GloFAS, TriggerCit, and reports from the Nepal Disaster Risk Reduction portal (NDRR)<sup>22</sup>. The relevance of the locations extracted with our approach is partial, but significant in some cases. The limited amount of data is possibly related to the specificity of our queries, which are focused on flood terms, while 50 out of 59 incidents reported by NDRR have a “Heavy Rainfall” or “Landslides” label. The impact of a tailored dictionary should be further investigated.



**Figure 6. Nepali case study - June 16-17, 2021**

### Assessment of the results


TriggerCit outputs an alert on June 16 due to the increase of tweets compared to the days before considered as baseline (Fig. 6). This trigger is related to a specific local event. Heavy rainfall caused a number of floods and landslides, across different regions of Nepal during the following weeks. Comparing the two periods, it appears that the peaks of tweets are more likely for sudden unexpected events, rather than for an expected event which is developing over time with many local events over an extended area. The corresponding maps and reference maps are shown in Fig. 7. The ability of TriggerCit of keeping track of these events is still under evaluation.

### COMPARISON WITH ALTERNATIVE SOURCES

UNOSAT Rapid Mapping Service<sup>23</sup> provides satellite image analysis during humanitarian emergencies, both natural disasters, and conflict situations. With a 24/7 year-round availability to process requests, a team of experienced analysts ensures timely delivery of satellite imagery-derived maps, reports, and data ready for direct inclusion in Geographic Information Systems (GIS) according to needs. The service can be activated by UN offices and agencies

<sup>22</sup>Events of flood, heavy rain and landslides reported from <http://drportal.gov.np/> for July 1-2

<sup>23</sup><https://www.unitar.org/maps/unosat-rapid-mapping-service>



**Figure 7. Nepali case study - July 1-2, 2021**

Government agencies, the Red Cross and Red Crescent Movement (ICRC and IFRC), international and regional organizations, and Humanitarian Non-Governmental Organizations (NGOs). As shown in Table 4, UNOSAT was activated on September 28, 2021 for the flood event in Thailand and on June 30, 2021 for the flood in Nepal. The activation data refers to the day when UNOSAT received the request to support the flood events. We were not able to find any flood activation from Copernicus EMS<sup>24</sup> in Thailand and Nepal respectively.

GDACS reported floods on September 27, 2021 for Thailand<sup>25</sup>, and on June 28, 2021 for Nepal<sup>26</sup> both with a green GDACS score<sup>27</sup> of 0.5. The event on June 16, 2021 was not recorded in GDACS following floods reported on the GDACS website<sup>28</sup>. We found the events reported by GloFAS using the GloFAS Map Viewer<sup>29</sup> on September, 24 in Thailand with a yellow and red score of Rapid Impact Assessment and on June 28, 2021 in Nepal with a orange score of Rapid Impact Assessment. FloodList also reported news of flood in Thailand<sup>30</sup> on September 27, 2021, and in Nepal on July 4, 2021<sup>31</sup> and on June 16, 2021<sup>32</sup>.

In the two scenarios, TriggerCit respectively returned a large scale and a local scale alert, both in a timely manner and accompanied by a valid geographical description, while providing information complementary to existing disaster alert mechanisms (e.g., images).

## LIMITATIONS AND FUTURE WORK

Linguistic limitations while mining multi-lingual tweets are mentioned in most of the works discussed in the state-of-the-art. The task is usually easier if the tweets are in English, given the quality of linguistic resources and the number of posts available. The use of native languages for communication has always posed a significant challenge to social media mining. In this paper, we propose a multi-language and extensible approach towards a solution to

<sup>24</sup><https://emergency.copernicus.eu/mapping/#zoom=2&lat=27.6533&lon=-25.0083&layers=0BT00>

<sup>25</sup><https://www.gdacs.org/report.aspx?eventtype=FL&eventid=1101113>

<sup>26</sup><https://www.gdacs.org/report.aspx?eventid=1100949&episodeid=1&eventtype=FL>

<sup>27</sup>[https://www.gdacs.org/Knowledge/models\\_FL.aspx](https://www.gdacs.org/Knowledge/models_FL.aspx)

<sup>28</sup><https://www.gdacs.org/flooddetection/currentfloods.aspx>

<sup>29</sup><https://www.globalfloods.eu/glofas-forecasting/>

<sup>30</sup><https://floodlist.com/asia/thailand-tropical-storm-dianmu-floods-september-2021>

<sup>31</sup><https://floodlist.com/asia/nepal-floods-landslides-july-2021>

<sup>32</sup><https://floodlist.com/asia/bhutan-and-nepal-flash-floods-leave-at-least-10-dead-many-missing>

Dates / Countries	Thailand	Nepal
UNOSAT activation	28/09/2021	30/06/2021
Copernicus EMS activation	None	None
GDACS Disaster Alerts	27/09/2021 (Green Alert)	28/06/2021 (Green Alert)
GloFAS	24/09/2021	28/06/2021
FloodList reported news	27/09/2021	04/07/2021
TriggerCit	26/09/2021	02/07/2021

**Table 4. Timeline of alerts for the Thailand and Nepal activations**

this problem. Further work is needed to develop fast and reliable seed dictionaries for social media monitoring with TriggerCit. For this purpose, several directions have to be investigated, leveraging on local resources, including local administrations and NGOs, and citizens. They can provide information about the nature of repeating local events, such as monsoons, and their consequences, to be mapped to keywords. Citizen scientists can also support in the evaluation of extracted images with crowdsourcing, as it has been proposed in the Crowd4EMS project (Ravi Shankar et al. 2019), as well as of the resulting keywords of the TriggerCit methodology (see Step 7). A software library could then be used to automate the aggregation of annotations (e.g., Cerquides and Mülâyim 2022).

As discussed in the paper, Twitter has different usage rates in different countries. Alternatives to Twitter include other social media like Facebook, Telegram, and WhatsApp. These are used to facilitate communication between volunteers and emergency services, enable sharing of information within a community, providing swift updates on emergency situations in real-time. However, even though such platforms contain a wealth of information, access to data is usually restricted (e.g., for privacy concerns). On the other hand, systems like Facebook Disaster Maps provide aggregated insights for crisis response and recovery and its use should be further investigated. In addition, as some tweets contain URLs pointing to other contents, such as web pages (e.g., newspapers) or posts in other social networks, these contents can also be crawled in order to obtain more information about the event. Moreover, the information extracted from one social media could be leveraged to crawl other platforms iteratively, as in Autelitano et al. 2019.

Further directions to be explored in order to enhance TriggerCit capabilities include:

- Training a supervised learning classifier for triggering  
We are still in the phase of collecting high-quality time series data related to a representative number of flood events. Once a large dataset is available, we plan to experiment with supervised learning approaches, in order to automate the triggering decision and validate it against a variety of events (e.g., different scales and areas). Such a classifier should also include contextual data about the event, information about flooded areas (such as the ones automatically classified in Nemni et al. 2020), and about the readiness of the areas in terms of preparation measures for recurring flood events.
- Automating dictionary management  
Our system needs high-quality dictionaries for each language. Part of the future work will focus on automatically deriving event-specific keywords, both in general and adapted to a particular event. This automation would enhance the sensing phase, in which the system is monitoring social media for potential triggers. On the other hand, online adaptation to a specific event could increase the recall of the process.

## CONCLUDING REMARKS

The paper proposed a multilingual methodology to provide early flood alerts with associated images and locations. The results in two case studies in Asia in 2021 show that the alert is consistent in the Thailand case and that a timely evidence could have been provided for the Nepal floods. Based on this analysis, TriggerCit could be deployed at UNOSAT into a cloud centralised service at CERN built on Kubeflow, a machine learning platform running on Kubernetes, allowing the integration with automatic flood detection from satellite data.

**Acknowledgements** The work at Politecnico di Milano, University of Geneva and IIIA-CSIC was funded by the European Commission H2020 Project Crowd4SDG, #872944.

The work at the United Nations Satellite Centre (UNOSAT) is part of the operational mapping service funded by the Norwegian Ministry of Foreign Affairs.

## REFERENCES

- Autelitano, A., Pernici, B., and Scalia, G. (July 2019). “Spatio-temporal mining of keywords for social media cross-social crawling of emergency events”. In: *GeoInformatica* 23.3, pp. 425–447.
- Cerquides, J. and Mülâyim, M. O. (Jan. 2022). *crowdanalysis: A software library to help analyze crowdsourcing results*. <https://doi.org/10.5281/zenodo.5898579>.
- Fohringer, J., Dransch, D., Kreibich, H., and Schröter, K. (2015). “Social media as an information source for rapid flood inundation mapping”. In: *Natural Hazards and Earth System Sciences* 15.12, pp. 2725–2738.
- Havas, C., Resch, B., Francalanci, C., Pernici, B., Scalia, G., Fernandez-Marquez, J. L., Van Achte, T., Zeug, G., Mondardini, M. R. R., Grandoni, D., et al. (2017). “E2mC: Improving emergency management service practice through social media and crowdsourcing analysis in near real time”. In: *Sensors* 17.12, p. 2766.
- HydroConference | WMO* (Accessed on 01/21/2022). <https://hydroconference.wmo.int/en/partners/products/global-flood-awareness-system-glofas>.
- Imran, M., Ofli, F., Caragea, D., and Torralba, A. (2020). “Using AI and Social Media Multimodal Content for Disaster Response and Management: Opportunities, Challenges, and Future Directions”. In: *Information Processing & Management* 57.5, p. 102261.
- Kankanamge, N., Yigitcanlar, T., Goonetilleke, A., and Kamruzzaman, M. (2020). “Determining disaster severity through social media analysis: Testing the methodology with South East Queensland Flood tweets”. In: *International journal of disaster risk reduction* 42, p. 101360.
- Kavallieros, D., Leventakis, G., Malliaros, S., Daniilidis, I., and Grizis, V. (2015). “PPDR Information Systems-A Current Status Review Report.” In: *ISCRAM*.
- Kirsch, B., Giesselbach, S., Knodt, D., and Rüping, S. (2018). “Robust End-User-Driven Social Media Monitoring for Law Enforcement and Emergency Monitoring”. In: *Community-Oriented Policing and Technological Innovations*, p. 29.
- Negri, V., Scuratti, D., Agresti, S., Rooein, D., Scalia, G., Shankar, A. R., Fernandez-Marquez, J. L., Carman, M. J., and Pernici, B. (2021). “Image-Based Social Sensing: Combining AI and the Crowd to Mine Policy-Adherence Indicators from Twitter”. In: *43rd IEEE/ACM International Conference on Software Engineering: Software Engineering in Society, ICSE (SEIS) 2021, Madrid, Spain, May 25-28, 2021*. IEEE, pp. 92–101.
- Nemni, E., Bullock, J., Belabbes, S., and Bromley, L. (2020). “Fully Convolutional Neural Network for Rapid Flood Segmentation in Synthetic Aperture Radar Imagery”. In: *Remote Sensing* 12.16.
- Olteanu, A., Castillo, C., Diaz, F., and Vieweg, S. (May 2014). “CrisisLex: A Lexicon for Collecting and Filtering Microblogged Communications in Crises”. In: *Proceedings of the International AAAI Conference on Web and Social Media* 8.1, pp. 376–385.
- Ravi Shankar, A., Fernandez-Marquez, J. L., Pernici, B., Scalia, G., Mondardini, M. R., and Di Marzo Serugendo, G. (2019). “Crowd4Ems: A crowdsourcing platform for gathering and geolocating social media content in disaster response”. In: *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 42, pp. 331–340.
- Al-Rfou, R., Kulkarni, V., Perozzi, B., and Skiena, S. (2015). “Polyglot-NER: Massive multilingual named entity recognition”. In: *Proceedings of the 2015 SIAM International Conference on Data Mining*. SIAM, pp. 586–594.
- Rivera, O., Calderon, J. F., Planchon, P., and Pernici, B. (Dec. 2021). “Evaluating the impact of floods on gender equality from social media evidence”. In: *2nd AISWN International Research Workshop on Women IS and Grand Challenges @ICIS*.
- Sakaki, T., Okazaki, M., and Matsuo, Y. (2013). “Tweet Analysis for Real-Time Event Detection and Earthquake Reporting System Development”. In: *IEEE Trans. Knowl. Data Eng.* 25.4, pp. 919–931.
- Scalia, G., Francalanci, C., and Pernici, B. (2021). “CIME: Context-aware geolocation of emergency-related posts”. In: *GeoInformatica*, pp. 1–33.
- Scotti, V., Giannini, M., and Cioffi, F. (2020). “Enhanced flood mapping using synthetic aperture radar (SAR) images, hydraulic modelling, and social media: A case study of Hurricane Harvey (Houston, TX)”. In: *Journal of Flood Risk Management* 13.4, e12647.
- Shoyama, K., Cui, Q., Hanashimaa, M., Sano, H., and Usuda, Y. (2021). “Emergency flood detection using multiple information sources: Integrated analysis of natural hazard monitoring and social media data”. In: *Science of the Total Environment* 767.144371, pp. 1–11.

- Singh, O. M., Padia, A., and Joshi, A. (Dec. 2019). “Named Entity Recognition for Nepali Language”. In: *2019 IEEE 5th International Conference on Collaboration and Internet Computing (CIC)*, pp. 184–190.
- Wang, R.-Q., Hu, Y., Zhou, Z., and Yang, K. (2020). “Tracking Flooding Phase Transitions and Establishing a Passive Hotline With AI-Enabled Social Media Data”. In: *IEEE Access* 8, pp. 103395–103404.
- Zajac, Z., Revilla-Romero, B., Aga Hirpa, F., Ntegeka, V., Salamon, P., Thielen, J., Burek, P., and Beck, H. (2014). “Recent advances in the Global Flood Awareness System (GloFAS)”. In: *EGU General Assembly Conference Abstracts*, p. 12182.