Landslide Detection Service Based on Composition of Physical and Social Information Services

Aibek Musaev, De Wang, Chien-An Cho, Calton Pu
Georgia Institute of Technology, Atlanta, Georgia
{*aibek.musaev, wang6, itszero, calton.pu*}@*gatech.edu*

***Abstract*—Social media have been used in the detection and management of natural hazards such as earthquakes. However, disasters often lead to other kinds of disasters, forming multi-hazards. Landslide is an illustrative example of a multi-hazard, which may be caused by earthquakes, rainfalls, water erosion, among other reasons. Detecting such multi-hazards is a significant challenge, since physical sensors designed for specific disasters are insufficient for multi-hazards. We describe LITMUS — a landslide detection service based on a multi-service composition approach that combines data from both physical and social information services by filtering and then joining the information flow from those services based on their spatiotemporal features. Our results show that with such approach LITMUS detects 25 out of 27 landslides reported by USGS in December and 40 more landslides unreported by USGS. Also, LITMUS provides a live demonstration that displays results on a web map.**

***Keywords*-landslide detection service; multi-service composition; social media; physical sensors; event detection**

I. Introduction

The detection of natural disasters is a significant and non-trivial problem. A traditional method relies on dedicated physical sensors to detect specific disasters, e.g., using seismometers for earthquakes. However, there are few physical sensors for the detection of multi-hazards such as landslides, which have multiple causes (earthquakes and rainfalls, among others) and happen in a chain of events. A more recent method explores the big data from social information services such as Twitter data streams functioning as social sensors [1]. High expectations have been placed on social sensors, since physical sensors (e.g., seismometers) are specialized for specific disasters. But despite some initial successes, social sensors have met serious limitations due to the big noise in big data generated by social sensors. For example, Twitter filter for the word “landslide” gets more tweets on the 70’s rock song “Landslide” [2] than landslide disasters that involve soil movement.

In this paper, we describe LITMUS — a landslide detection service that is based on a multi-service composition approach to the detection of landslides, a representative multi-hazard. Instead of trying to refine the precision and recall of event detection in each one of the physical and social information services[1], LITMUS composes information from a variety of sensor networks. The information services include both physical sensors (e.g., seismometers for earthquakes and weather satellites for rainfalls) and social sensors (e.g., Twitter and YouTube). Instead of trying to optimize the filtering process for each social sensor in isolation, LITMUS uses state-of-art filters for each social sensor, and then adopts geo-tagging to integrate the reported events from all physical and social sensors that refer to the same geo-location. Our work shows that with such integration LITMUS achieves better landslide detection when compared to an authoritative source.

LITMUS is one of the few works that makes use of the composition of multiple heterogeneous information services. It is not tied to disaster detection and can be applied to other application areas involving service composition. This work presents a generic approach to the problem of composition of multiple heterogeneous information services and uses landslide detection as an illustrative example. Traditional approach to the composition of web services makes strong assumptions about services, which it then uses to select services when composing a new service, such as quality of service [3] or service license compatibility [4]. In practice, the real world services do not satisfy such assumptions. The assumption we make in our work is that more information services should provide a more solid result and we demonstrate that it is the case with LITMUS.

The rest of the paper is organized as follows. Section II provides a brief overview of the landslide detection service and its components, Section III introduces the supported physical and social information services, and Section IV describes implementation details of each system component. In Section V we present an evaluation of landslide detection using real data and compare the results generated by LITMUS with an authoritative source. We summarize related work in Section VII and conclude the paper in Section VIII.

II. Framework Overview

A. System Requirements

Web service LITMUS was designed to detect landslides based on a multi-service composition approach that com-

[1]Throughout this paper the terms ”information service” and ”sensor” will be used interchangeably.

bines data from physical and social information services. Physical services should include earthquake and rainfall real-time feeds as possible causes of landslides. LITMUS should also support various social information services, which we expect to help detect landslides. However, the data from the social services must be filtered as they often contain a lot of noise. The system should also adopt geo-tagging to integrate the reported events from all physical and social sensors that refer to the same geo-location. Finally, a web client was designed to demonstrate LITMUS functionality by displaying detected landslides on a web map.

B. Implementation Overview

Based on the requirements described above, we implemented 3 independent components that perform filtering, integration and semantics-aware detection shown in Figure 1. The Filtering component downloads the data from social and physical sensors and filters out noise from social sensors. The Integration component combines the filtered data from social sensors with the data from physical sensors based on a Bayesian model integration strategy to generate a list of potential landslide locations. The last component performs semantics-aware detection of landslides by grouping locations related to the same event and excluding the results that are not current.

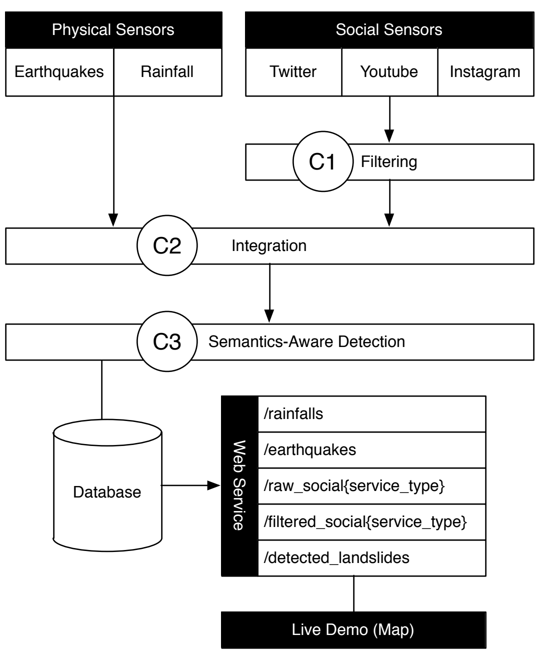

Figure 1. Overview of LITMUS Framework

LITMUS provides access to its resources via a web service. This web service is implemented in a representational state transfer (REST) style [5]. The architectural properties of this style are driven by several constraints, including client-server architecture and stateless communication. Client-server architecture ensures separation of concerns between clients and servers. This means, for example, that clients are not concerned with data storage, which is handled by servers. Servers are not concerned with the user interface or user state. The client-server interaction is further constrained by stateless communication, which means that no client context is stored on the server and that each client request contains all necessary information to be executed by the server. In addition to the described constraints, the central principle in REST is support for resources that are sources of information provided by the web service. Each resource is referenced with a global identifier, such as URI. In our landslide detection service the resources are the physical sensor feeds, the raw social feeds that are downloaded, the filtered social feeds that are processed by the system and the resulting feed of detected landslides as shown in Figure 1 in the web service component.

We implemented a web service demonstration consuming the resources provided by LITMUS, which is located in the GRAIT-DM portal accessible at [6]. GRAIT-DM is a SAVI project for Global Research on Applying Information Technology to support Effective Disaster Management.

III. Physical and Social Information Services

The physical information services supported by LITMUS do not provide information about landslides directly, but they provide reports about other kinds of disasters, which may be possible causes of landslides. In particular, we use a real-time earthquake activity feed from the US Geological Survey (USGS) agency [7]. This feed is updated every minute and provides information about earthquakes of various magnitute. For this paper we collect data on earthquakes of 2.5 magnitude and higher. The data is provided in a convenient GeoJSON format, which among other things provides time, magnitude, latitude, longitude, name of the place and ID, which is used to avoid duplicate records in the system.

Another physical information service supported by LITMUS is provided by the Tropical Rainfall Measuring Mission (TRMM) [8], which is a joint project between NASA and the Japan Aerospace Exploration Agency (JAXA). This project generates reports based on the satellite data of the areas on the planet that have experienced rainfalls within the past one, three and seven days. The reports are provided in multiple formats, including reports in a uniform manner on the project's web page, from which we parse and extract the rainfalls data.

LITMUS also supports various social information services, which we expect to help detect landslides, since there are no physical sensors that would detect landslides directly. Currently we support Twitter as an example of a text based social network, YouTube as an example of a video based social network and Instagram as an example of an image based social network. Each of these social information services is among the leading social networks in their respective areas – 500 million tweets are posted per day[2], 55 million photos are sent per day[3] and 100 hours of videos are uploaded per minute[4].

Next we will present the implementation details of each system component.

IV. System Components

C1. Filtering Component

The C1. Filtering Component applies only to social information services. This is due to the fact that all information regarding both earthquakes and rainfalls is considered relevant for landslide detection purposes. Also, physical information services provide data with geo-coordinates, hence there is no need to apply the geo-tagging processing either.

To remove noise from social sensors, we process the downloaded data in a series of filtering steps. There are four stages in this process that filter out items, which are neither related to landslides (steps F1, F2 and F4) nor are useful to LITMUS due to lack of geo-location (step F3).

F1. Filter based on search terms: Each social network provides search API based on keywords for software developers. The system periodically downloads the data from each social sensor based on "landslide" and "mudslide" keywords. The period is currently set to 30 minutes, which can be modified if necessary.

F2. Filter based on stop words & phrases: The social information services require additional filtering as they contain a lot of items unrelated to landslides and most of the time they are not geo-tagged either. The following is a set of frequent examples of unrelated items from the social information services:

- "Landslide" song by Stevie Nicks from Fleetwood Mac: "Climbed a mountain and I turned around, and I saw my reflection in the snow-covered hills, and the landslide brought me down. -FleetwoodMac"
- Used as an adjective describing an overwhelming majority of votes or victory: "Robert Mugabe's party claims landslide victory in Zimbabwe's key election as ... - Daily Mail http://t.co/Hf4sVU3E8F"
- Lyrics from "Bohemian Rhapsody" by Queen: "Caught in a landslide, no escape from reality..."

The first two items can be filtered out using a simple exclusion rule based on the presence of stop words "FleetwoodMac" and "election". The third item is filtered out using another exclusion rule based on the presence of stop phrases that currently include the lyrics of some popular songs, e.g. "no escape from reality".

But other unrelated to landslide items require a more sophisticated algorithm, including filtering based on geo-location and filtering based on penalized classification, described next.

[2]https://about.twitter.com/company
[3]http://instagram.com/press/
[4]http://www.youtube.com/yt/press/

F3. Filter based on geo-location: To detect landslides within a particular period, we need to determine their locations. The data from the physical sensors already contains geo-coordinates. However, the data from the social sensors is usually not geo-tagged although each social network provides support for users to disclose their location. So, if an item has not been geo-tagged already, then we suggest to look for mentions of place names that refer to locations of landslides in the item's text.

For details of the geo-tagging algorithm used in LITMUS we refer the reader to [9]. Next we will describe the changes that we made in this algorithm.

The geo-tagging algorithm extracts geographical terms (geo-terms) from incoming messages and assigns geo-coordinates to each geo-term using a gazetteer, which is a dictionary that maps places to geo-coordinates. In our system we use a public gazetteer data from the GeoNames database that covers all countries and contains over 10 million places [10]. Our geo-tagging algorithm uses a subset of this data, namely 448k places. This subset includes countries, administrative divisions of countries of first to fourth orders, cities with population greater than 1,000 people and islands. In the future we plan to increase this subset by including even more detailed places.

Some news sources mentioned in social media data, such as "Boston Globe" or "Jamaica Observer", contain valid geo-terms which must be removed from consideration by the geo-tagging algorithm; otherwise it would return incorrect results. That is why LITMUS maintains a list of major news sources, including "Boston Globe" and "Jamaica Observer". Consider the following tweet: "Boston Globe - Typhoon, mudslides kill 14 in Japan; 50 missing http://t.co/nEUbk60Pzl." "Boston" is positioned closer to the landslide keyword "mudslides" than "Japan", however it is a part of the "Boston Globe" news source that LITMUS automatically removes from consideration, such that the correctly extracted geo-term is "Japan".

F4. Filter based on classification: Majority of items returned by social sensors are not relevant to landslides, even though they contain landslide keywords and valid geo-terms. The following are examples of irrelevant items with respect to landslide disasters that contain valid geo-terms:

- Extracted geo-term "California": "Laying on a landslide with a bag of California to smoke."
- Extracted geo-term "New York City": "Lupica: As Bill de Blasio takes the mayoralty of New York City, let's not forget Chris Christie's landslide vi http://t.co/0CoR1zQY55"

To filter out such irrelevant items LITMUS employs machine learning binary classification, which automatically labels each item as either relevant or irrelevant based on a classifier model built from a training set containing labeled items. This method, however, generates both correct and incorrect labels. For successful labeling we propose to convert the filtering problem of each item to the filtering problem of landslide locations, i.e. geo-terms. In most cases a particular geo-term is mentioned in multiple incoming items from social sensors. Each social item is labeled by a machine learning classifier, so there are multiple classification results for each geo-term. Although this method generates both correct and incorrect labels for a particular geo-term, but on average it provides more correct than incorrect labels. The idea of penalized classification uses this heuristics to improve the results of classification by accepting the label assigned to the majority of items for each geo-term and only considering locations whose majority label is positive.

C2. Integration Component

To generate a list of potential landslide locations, LITMUS combines the data from physical information services with filtered and geo-tagged data from social information services. This is a two stage process, where a grid-based landslide location estimation is followed by an integration of results from multiple services.

To estimate landslide locations, we propose to represent the surface of the Earth as a grid of cells. Each geo-tagged item is mapped to a cell in this grid based on the item's coordinates. After all items are mapped to cells in this grid, the items in each non-empty cell are used for computing an integrated landslide score. The size of the cells is equal to 2.5 minutes in both latitude and longitude, which corresponds to the resolution of the Global Landslide Hazard Distribution [11] that we plan to add as an additional physical sensor to the system. This is the maximum resolution supported in LITMUS. The actual resolution is driven by the precision of the geo-tagging algorithm described earlier.

After mapping the items from each sensor to cells in this grid, which represent potential landslide locations, we calculate the probability of landslide occurrence in location cells based on a Bayesian model integration strategy. Here is the description of this strategy. We use a subscript i to distinguish different sensors. Suppose we have a cell x and ω is the class associated with x, either being true or false. Then, assuming a hidden variable Z for an event to select one sensor, a probability for a class ω given $x, P(\omega|x)$, can be expressed as a marginal probability of a joint probability of Z and ω:

$$P(\omega|x) = \sum_i P(\omega, Z_i|x) = \sum_i P(\omega|Z_i, x) P(Z_i|x) \tag{1}$$

Here, we use external knowledge $P(Z_i|x)$ to express each sensor's confidence given x. For instance, if a certain sensor becomes unavailable, the corresponding $P(Z_i|x)$ will be zero. Also one could assign a large probability for the corresponding $P(Z_i|x)$ if one sensor dominates over other sensors.

In our experiment, we use prior F-measure R from the August data as the confidence for each sensor since F-measure provides a balance between precision and recall, namely $F\text{-}measure = 2 * \frac{precision*recall}{precision+recall}$. We use the August data in our experiments, because there were many landslides reported by USGS during this month, so the data collected during this month is representative. To generate the results in the range from 0 to 1, we normalize the values of F-measure into a scale between 0 and 1 first. Taking all items from each sensor into account, the formula will be further converted into the following format:

$$P(\omega|x) = \sum_i R_i \frac{\sum_j POS_{ij}^x - \sum_j NEG_{ij}^x - \sum_j STOP_{ij}^x}{\sum_i N_i^x} \tag{2}$$

Here, R_i denotes the normalized prior F-measure of sensor i from historic data. POS_{ij}^x denotes positively classified items from sensor i in cell x, NEG_{ij}^x denotes negatively classified items from sensor i in cell x, $STOP_{ij}^x$ denotes the items from sensor i in cell x that have been filtered out using stop words and stop phrases, and N_i^x is a total number of items from sensor i in cell x.

C3. Semantics-Aware Detection Component

The actual number of landslides is usually less than a total number of potential landslide locations returned by the Integration component. Some of the potential landslide locations may be referring to the same event. For example, this may happen when users refer to the same event by using geo-terms of different level of detail. Consider the following tweets that describe a landslide, which occurred in Italy in December:

- Extracted geo-term "italy": "Giant craters were ripped out of roads in Italy and homes and shops sank into the ground after a major landslide: http://t.co/JZ6l63vXLH"
- Extracted geo-term "montescaglioso": "VIDEO: Landslide rips apart Italy roads: Heavy rains and floods cause a powerful landslide in the southern Italian town of Montescaglioso."

Both geo-terms extracted from these tweets are valid, but they describe the same event, which should be correctly detected by the Semantics-Aware Detection component. Another possible scenario of having multiple geo-terms describing the same event can be described by the following tweets:

- Extracted geo-term "wadhurst": "Christmas come early thanks to #southeastern, the bad santa of train services. Landslide at Wadhurst has block line. Working from home."
- Extracted geo-term "hastings": "Train delayed due to landslide. That's a first for the Hastings line."

Social Sensors	F1.Filter based on search terms	F2.Filter based on stop words & phrases	F3.Filter based on geo-location	F4.Filter based on classification
Twitter	45391	24993	7879	2815
Instagram	1418	1263	308	15
YouTube	4890	3936	557	318

Table I
OVERVIEW OF FILTERING RESULTS

Metrics	F1.Filter based on search terms	F2.Filter based on stop words & phrases	F3.Filter based on geo-location	F4.Filter based on classification
Signal-to-Noise Ratio	0.05	0.10	0.44	1.90
Information Gain	-	0.028	0.121	0.197

Table II
EVALUATION RESULTS OF FILTERING COMPONENT

These two landslide locations are actually related as shown in the following tweet:

- Extracted geo-terms "wadhurst" and "hastings": "Avoid the trains on the Hastings line folks. Word is there's been a landslide on the line near #Wadhurst #UKstorm"

Semantics-Aware Detection component must also handle temporal issues by excluding results that refer to the past or future events:

- "The Kedarnath disaster in Uttarakhand, India in June remains the worst landslide accident of 2013 to date http://t.co/Mf31ztjwQ2"

Even though the year matched the year of the evaluation period, the month did not, hence the landslide locations extracted from this message must be excluded from the final result.

The Semantics-Aware Detection component is currently semi-automated. LITMUS is able to group landslide locations that were referred to in the same message and also to exclude messages containing references to either past or future years. We plan to improve the performance of this component as part of future work.

V. EVALUATION USING REAL DATA

To evaluate the performance of our landslide detection service we designed three sets of experiments. We start with evaluating the performance of the filtering process of social information services. Next we compare the effectiveness of three multi-service composition strategies for landslide detection. The final experiment provides the detection comparison results between LITMUS and an authoritative source.

A. *Evaluation of Filtering Component*

In this experiment, we view different social media as our social sensors and process the data from these sensors in a series of filtering steps. For simplicity, we only focus on the textual description of each item during filtering. For Instagram, the textual description is an image's caption text. For YouTube, the textual description is a concatenation of a video's title and description. And for Twitter, the textual description is the text of a tweet itself. We use December 2013 as our evaluation period. Table I shows the total number of items downloaded during this month. For each filtering step, Signal and Total indicate the remaining items after the filtering steps, where Signal is a number of items from social sensors that are relevant to landslide detection and Total is a total number of items from social sensors.

There are two metrics that we use to evaluate the performance of the filtering steps in Table II, namely Signal-to-Noise Ratio (SNR), and Information Gain (IG), where for each filtering step i:

$$SNR_i = \frac{Signal_i}{Noise_i} = \frac{Signal_i}{Total_i - Signal_i} \tag{3}$$

Here, $Signal_i$ is the number of relevant items remaining in step i, $Total_i$ is the number of all items remaining in step i.

$$IG_i(T_{i-1}, a_i) = H(T_{i-1}) - H(T_{i-1}|a_i) \tag{4}$$

Here, we consider the filtering process as a binary classification problem which has two classes: relevant and irrelevant. T_{i-1} denotes a set of training examples before step i. a_i is the attribute (filtering condition) we used in step i. H denotes information entropy. The filtering conditions in the process are considered as attributes. For instance, the F2 filter based on stop words and phrases will be converted into a classifier based on a boolean attribute whether an item contains stop words and phrases. Information gain measures the relevance of attributes.

It can be seen from Table II that SNR is improving after each filtering step and eventually exceeds 1. Information gain ranks the relevance of the filtering conditions, which shows the rank of filters in decreasing order: filter based on classification, filter based on geo-location, and filter based on stop words and phrases. For F1, we cannot calculate the values of Information Gain, since we are not able to know how many items there are in filtered out data and what the prior information entropy is.

B. Evaluation of Integration Component

For comparison purposes, we introduce two major baselines - OR and AND integration strategies and compare them with the proposed Bayesian model strategy. We have five sensors in total, including social sensors (Twitter, Instagram, and YouTube) and physical sensors (earthquakes and rainfalls). Considering that each sensor has one vote to a particular cell, the cell will obtain one or zero votes from each sensor. For OR integration strategy, we grant equal weight to five sensors. And we obtain the decision (whether a landslide happened or not) by combining the votes using boolean operation OR among five sensors. For social and physical AND integration strategy, we use boolean operation OR to combine the votes from social sensors and physical sensors separately first. And then we calculate the combined result by applying boolean operation AND between votes from social and physical sensors. For instance, if the votes from five sensors (Twitter, Instagram, YouTube, earthquakes, and rainfalls) are 1,1,0,0, and 0, the OR strategy will give a 1 in the end, but the social AND physical strategy will give a 0 in the integrated score.

We present the results of comparison between integration strategies in Figure 2. This figure shows that the Bayesian model strategy has 71% precision, 82% recall and 77% F-measure. The OR strategy has the highest recall at 100%, but also the lowest precision at 2%. The AND strategy shows improvement compared to the OR strategy as its F-measure is higher, but the Bayesian strategy shows the best performance overall.

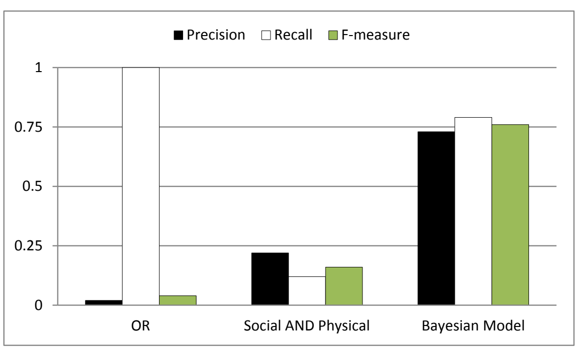

Figure 2. Comparison of Integration Strategies

C. Evaluation of Semantics-Aware Detection Component

In addition to a real-time seismic feed, USGS also presents a continually updated list of landslides as reported by other reputable sources, including Dailymail.co.uk, GlobalPost, HeraldNet, and Xinhuanet. In this experiment we compare the landslides provided by this authoritative source in December [12] versus the landslides detected by LITMUS during the same month — see Figure 3 for the results of this comparison. LITMUS was able to detect 25 out of 27 landslides reported by USGS. The 2 missed events were in Buncombe County, NC and Watchung, NJ. Both events affected very small areas, which is the reason why they did not attract significant public interests in social media and were missed by the system. In addition to the overlapping events detected by both LITMUS and USGS, our system managed to find 40 more landslide locations in December that were unreported by USGS. This is due to the fact that the USGS results are based on news sources, which can only report a limited amount of information. Whereas LITMUS employs multiple information services, plus the landslide detection process is automated, so we can reasonably claim the comprehensiveness of its results.

Figure 3. Landslide Detection by LITMUS vs USGS

VI. Web Service Demonstration

We developed a live demonstration [6] that consumes the resources provided by the web service. This web application shows live feeds from each resource described in the paper. The data from all feeds is displayed on a Google Map, which can be set to either Map or Satellite view by a user. A user can view detailed information about items from each feed — see Figure 4 for a detailed view of an item from a YouTube feed.

A separate feed shows a list of detected landslides that are a result of the multi-service analysis based on the Bayesian model integration strategy. A user can also view detailed information about all items that were used to make a decision regarding a landslide in each location — see Figure 5.

VII. Related Work

Many researchers have explored the use of social media to detect events and provide assistance, including TED

Figure 4. Example of video from the YouTube feed

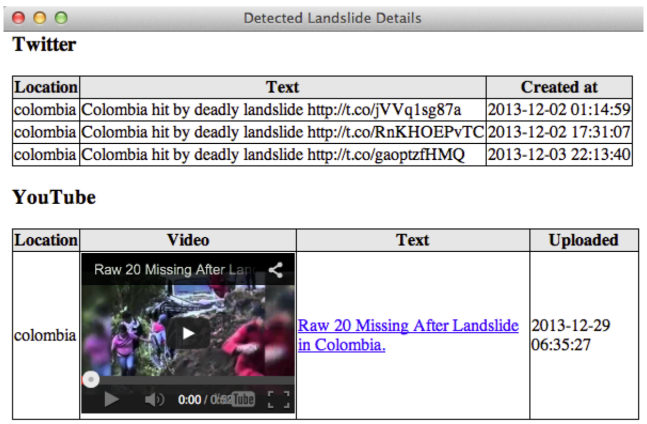
Detected Landslide Details

Twitter

Location	Text	Created at
colombia	Colombia hit by deadly landslide http://t.co/jVVq1sg87a	2013-12-02 01:14:59
colombia	Colombia hit by deadly landslide http://t.co/RnKHOEPvTC	2013-12-02 17:31:07
colombia	Colombia hit by deadly landslide http://t.co/gaoptzfHMQ	2013-12-03 22:13:40

YouTube

Location	Video	Text	Uploaded
colombia	Raw 20 Missing After Lan	Raw 20 Missing After Landslide in Colombia.	2013-12-29 06:35:27

Figure 5. Example details of a detected landslide

project [13], and social media analysis [1], [14]. Supporting techniques for these systems include topic detection [15], event detection [16], and alerts [17]. Besides Twitter, some studies targeted search engines, news, and blogs by analyzing spatiotemporal patterns [18], [19], [20]. The previous work on the integration of multiple data sources has been largely application-specific, for example, biostatistics [21], bioinformatics [22], and geographic information systems [23]. Supporting techniques include location estimation [24], and text categorization [27], [28], [29]. Our work continues development of LITMUS [9] by creating a landslide information service based on a multi-service composition approach that combines data from both physical and social information services.

Multi-service composition requires systems that can manage complex, heterogeneous, dynamic and distributed data, which may be quite large. Milanovic et al. [30] provided a survey of existing proposals for web service composition. Constantinescu et al. [31] presented an algorithm that supports dynamic service composition based on partial matches of input/output types. Truong et al. [32] proposed information quality metrics for identifying and reducing irrelevant information about web services. Service composition in LITMUS is static rather than dynamic, because the data from all of our sensor information services is downloaded at each cycle. As the number of information services supported by LITMUS grows, we plan to add support for dynamic service composition and execution.

Another important aspect for disaster detection systems is situational awareness. The challenge for social sensors is that users may use alias or location names in different granularities in messages resulting in inaccurate location information. Multiple studies have been done on location estimation for information from social networks based on content of tweets, e.g. [24]. [33] demonstrated a rapid unsupervised extraction of locations references from tweets using an indexed gazetteer. Our system also employs a public gazetteer and adopts a grid-based approach with customizable granularities in location estimation.

VIII. Conclusion

Multi-hazards are disasters with causally chained events such as the 2011 Tohoku earthquake triggering tsunami, which caused the Fukushima nuclear disaster, and landslides, often caused by earthquakes or rainfalls. Detecting such multi-hazards is a significant challenge, since physical sensors designed for specific disasters are insufficient for multi-hazards. As promising alternatives [1], social information services have difficulties with filtering the big noise in the big data being generated. We show that multi-service composition approach, which combines data from both physical and social information services, can improve the precision and accuracy of multi-hazard detection when the participating sensors are relatively independent of each other.

Applying this approach, we built a landslide detection service called LITMUS, which composes physical information services (USGS seismometers and TRMM satellite) and social information services (Twitter, Instagram, and YouTube). LITMUS provides a REST API for obtaining its resources, including social and physical sensor feeds and a list of detected landslides. A live demonstration [6] is developed that consumes these resources to display the results on a Google Map.

The effectiveness of landslide detection is evaluated using real world data collected in December 2013. Individual filtering results for each social sensor are provided followed by the full integration of 5 sensors applying a modified Bayesian model integration strategy that achieved 71% in precision, 82% in recall and 77% in F-measure for landslide detection, which is significantly better than the baseline integration strategies. A comparison is performed against an authoritative list compiled by USGS, which shows that LITMUS detects 25 out of 27 reported events as well as 40 more events unreported by USGS in December.

Finally, the coverage of landslides detected by LITMUS can be improved by supporting other languages in addition to English. We are also interested in detecting other kinds of events using LITMUS infrastructure, for example food poisoning and flu. Furthermore, our future work will involve an analysis of the image and video content as they contain more information in addition to their textual representation.

IX. Acknowledgements

This research has been partially funded by National Science Foundation by CNS/SAVI (1250260, 1402266), IUCRC/FRP (1127904), CISE/CNS (1138666), NetSE (0905493) programs, and gifts, grants, or contracts from Singapore Government, and Georgia Tech Foundation through the John P. Imlay, Jr. Chair endowment. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or other funding agencies and companies mentioned above.

References

[1] T. Sakaki *et al.*, “Earthquake shakes twitter users: real-time event detection by social sensors,” in *WWW*, 2010.

[2] S. Nicks, “Landslide (song),” http://en.wikipedia.org/wiki/Landslide_(song), accessed on 1/1/2014.

[3] S. Ran, “A model for web services discovery with QoS,” in *ACM SIGecom Exchanges*, 2003, vol. 4, no. 1.

[4] G. Gangadharan *et al.*, “Service license composition and compatibility analysis,” in *ICSOC*, 2007.

[5] R. T. Fielding *et al.*, “Principled design of the modern Web architecture,” *ACM Transactions on Internet Technology*, vol. 2, no. 2, 2002.

[6] GRAIT-DM, “Global Research on Applying Information Technology to support effective Disaster Management,” https://grait-dm.gatech.edu/demo-multi-source-integration/, accessed on 1/1/2014.

[7] USGS, “United States Geological Survey agency: Earthquake activity feed from the United States Geological Survey agency,” http://earthquake.usgs.gov/earthquakes/, accessed on 1/1/2014.

[8] TRMM, “Tropical Rainfall Measuring Mission: Satellite monitoring of the intensity of rainfalls in the tropical and subtropical regions,” http://trmm.gsfc.nasa.gov/, accessed on 1/1/2014.

[9] A. Musaev *et al.*, “LITMUS: Landslide detection by integrating multiple sources,” in *ISCRAM*, 2014.

[10] GeoNames, “GeoNames geographical database,” http://www.geonames.org/, accessed on 1/1/2014.

[11] CHRR et al., “Global Landslide Hazard Distribution,” http://sedac.ciesin.columbia.edu/data/set/ndh-landslide-hazard-distribution/, accessed on 1/1/2014.

[12] USGS, “United States Geological Survey agency: Recent landslide events,” http://landslides.usgs.gov/recent/index.php?year=2013&month=Dec, accessed on 1/1/2014.

[13] M. Guy *et al.*, “Integration and dissemination of citizen reported and seismically derived earthquake information via social network technologies,” vol. 6065, 2010.

[14] M. Kitsuregawa *et al.*, “Analytics for info-plosion including information diffusion studies for the 3.11 disaster,” in *Web-Age Information Management*, 2011, vol. 6897.

[15] X. Wang *et al.*, “Real time event detection in twitter,” in *Web-Age Information Management*, 2013, vol. 7923.

[16] T. Hua *et al.*, “STED: semi-supervised targeted-interest event detectionin in twitter,” in *KDD*, 2013.

[17] S. Kumar *et al.*, “Whom should i follow?: identifying relevant users during crises,” in *HT*, 2013.

[18] L. Backstrom *et al.*, “Spatial variation in search engine queries,” in *WWW*, 2008.

[19] Q. Mei *et al.*, “A probabilistic approach to spatiotemporal theme pattern mining on weblogs,” in *WWW*, 2006.

[20] K. Radinsky *et al.*, “Mining the web to predict future events,” in *WSDM*, 2013.

[21] N. J. Horton *et al.*, “Regression analysis of multiple source and multiple informant data from complex survey samples,” *Statistics in Medicine*, vol. 23, 2004.

[22] I. Huopaniemi *et al.*, “Multivariate multi-way analysis of multi-source data,” *Bioinformatics*, vol. 26, no. 12, 2010.

[23] P. H. Swain *et al.*, “Multisource data analysis in remote sensing and geographic information processing,” Purdue University, Tech. Rep., 1985. [Online]. Available: http://docs.lib.purdue.edu/larstech/78/

[24] Z. Cheng *et al.*, “You are where you tweet: a content-based approach to geo-locating twitter users,” in *CIKM*, 2010.

[25] E. Lex *et al.*, “Efficient cross-domain classification of weblogs,” *International Journal of Intelligent Computing Research*, vol. 1, no. 1, 2010.

[26] S. J. Pan *et al.*, “Cross-domain sentiment classification via spectral feature alignment,” in *WWW*, 2010.

[27] Y. Zhen *et al.*, “Cross-domain knowledge transfer using semi-supervised classification,” in *AI 2008: Advances in Artificial Intelligence*, 2008, vol. 5360.

[28] N. Milanovic *et al.*, “Current solutions for web service composition,” *IEEE Internet Computing*, vol. 8, no. 6, 2004.

[29] I. Constantinescu *et al.*, “Large scale, type-compatible service composition,” in *ICWS*, 2004.

[30] H. L. Truong *et al.*, “On identifying and reducing irrelevant information in service composition and execution,” in *WISE*, 2010.

[31] E. A. Sultanik *et al.*, “Rapid geotagging and disambiguation of social media text via an indexed gazetteer,” in *ISCRAM*, 2012.