



Improving remote sensing flood assessment using volunteered geographical data

E. Schnebele and G. Cervone

Dept. of Geography and Geoinformation Science, George Mason University, 4400 University Drive, Fairfax, VA, USA

Correspondence to: E. Schnebele (eschnebe@gmu.edu) and G. Cervone (gcervone@gmu.edu)

Received: 6 September 2012 – Published in Nat. Hazards Earth Syst. Sci. Discuss.: –
Revised: 19 November 2012 – Accepted: 23 January 2013 – Published: 19 March 2013

Abstract. A new methodology for the generation of flood hazard maps is presented fusing remote sensing and volunteered geographical data. Water pixels are identified utilizing a machine learning classification of two Landsat remote sensing scenes, acquired before and during the flooding event as well as a digital elevation model paired with river gage data. A statistical model computes the probability of flooded areas as a function of the number of adjacent pixels classified as water. Volunteered data obtained through Google news, videos and photos are added to modify the contour regions. It is shown that even a small amount of volunteered ground data can dramatically improve results.

flooding (Sanyal and Lu, 2004). The use of Landsat data for flood assessment can be highly effective. Frazier and Page (2000) employed a supervised maximum-likelihood classification to map water bodies with Landsat Thematic Mapper (TM), with an overall accuracy of over 97%. Although effective for detecting water, the satellite's orbit revisit time can constrain data availability making it difficult to create a comprehensive time series of a flood event. Cloud and vegetative cover can obscure surface measurements when utilizing optical data, often resulting in partial coverage and incomplete flood assessment.

Numerous attempts have been proposed to overcome the limitations of remote sensing data, often by supplementing them with additional data to provide a more accurate and comprehensive flood assessment. Laura et al. (1990), Townsend and Walsh (1998) have proposed the use of RADAR remote sensing data for the assessment of floods. RADAR has the unique advantage of penetrating through canopy and clouds, and can easily distinguish water bodies from most other land cover types. However, RADAR data is not widely available, and usually have limited swaths with long revisit times. Efforts have been made toward increasing RADAR's availability and accessibility. Wang et al. (2002) have proposed the integration of Landsat TM data with a digital elevation model (DEM) and river gage data to create a comprehensive assessment of flood depth under forest and cloud canopy. Although river gage data is usually sparse and not universally available, especially in more remote areas; this methodology proved very robust and is routinely used for flood assessment. The research described in this paper is inspired by Wang's methodology, where Landsat, DEM, and river gage data are used collectively in an attempt to improve flood analysis. The main difference consists in the

1 Introduction

The ability to produce accurate and timely flood assessments is a critical safety tool for flood mitigation and response. Several methodologies have been developed to assess the risks associated with flooding by using ground measurements such as precipitation, water flow or level (e.g. Richter et al., 1998; Apel et al., 2006). Satellite remote sensing data have been utilized for flood assessment because of their high spatial resolution and capacity to provide information for areas of poor accessibility or lacking in ground measurements (Smith, 1997). High resolution satellite data is particularly useful for the spatial analysis of water pixels. When data before and after a flood event are available, it is possible to classify land cover change, and thus identify which areas are flooded.

The Landsat satellite program has been collecting data about the Earth and its environment since the 1970s, and has been employed to monitor and mitigate the impacts of

fusing methodology employed in this study to integrate the different data sources.

An emerging and quickly growing data source not yet fully utilized with respect to natural hazards is volunteered geographic information (VGI) (Goodchild, 2007). This general class of data, voluntarily contributed and made available, contain temporal and spatial geographical information. Data sources include pictures, videos, sounds, text messages, etc. Due to the spread of the internet to mobile devices, an unprecedented and massive amount of ground data have become available, often in real-time. Some data are geolocated automatically, while others can be geolocated by analyzing content.

Although volunteered data is often published without scientific intent, and usually carry little scientific merit, it is still possible to mine mission critical information. For example, during hurricane Katrina, geolocated pictures and videos searchable through Google provided early emergency response with ground-view information. These data have been used during major events, with the capture in near real-time the evolution and impact of major hazards (De Longueville et al., 2009; Pultar et al., 2009; Heverin and Zach, 2010; Vieweg et al., 2010; Acar and Muraki, 2011; Verma et al., 2011; Earle et al., 2012; Tyshchuk et al., 2012).

This work is based on a specific subset of this general class of data, namely photos, videos, and news. Volunteered photos about natural hazards have emerged as a data source during crisis and hazardous events. Liu et al. (2008); Hyvärinen and Saltikoff (2010); McDougall (2011); Zhang et al. (2012) show how photos from Flickr have been used to derive local meteorological information, capture and record the physical features of an event, and identify and document flood height.

Recently, volunteered data have been evaluated for estimating flood inundation depth and for mapping flood extent (Poser and Dransch, 2010; McDougall, 2011). These potentially valuable, real-time data have yet to be regularly applied in large scale disaster relief situations for multiple reasons, including difficulties of authentication and confirmation, questions of quality and reliability, and difficulties associated with harvesting data from heterogeneous and non-structured sources (Flanagin and Metzger, 2008; Schlieder and Yanenko, 2010; Tapia et al., 2011).

This paper proposes a new methodology that leverages data freely obtained from the internet to improve flood hazard estimation. Combining the high spatial resolution and reliability of satellite imagery with the high temporal resolution of ground data takes advantage of the strengths of both data types while allowing for mutual data confirmation.

Despite the non-scientific nature of volunteered information, the integration of these data with traditional data sources offers an opportunity to include new and additional information in flood extent mapping. It is assumed that ground truth data is not available, and therefore

the quantitative analysis of the results discusses the changes introduced by fusing the different data layers. The novelty of this study is the development of a methodology that takes advantage of “citizens as sensors”, as discussed in Goodchild (2007), to fuse observations culled from social media with satellite and topographic data for flood assessment.

A case study is presented for the May 2011 flooding of the Mississippi River. This was one of the worst floods since the Great Flood of 1927. In Memphis, TN the Mississippi River crested at 14.6 m, the highest crest since 1937, which caused the evacuation of approximately 1300 homes. The methodology was implemented using the R statistical package¹.

2 Data

2.1 Volunteered data

The data used in this study have been downloaded using the Google search engine through their photos, videos and news portal. They included sources from Flickr, YouTube, Weather Underground, Wikipedia, and abc24.com. In particular, videos ($n = 6$) and photos ($n = 8$) from the first two weeks of May 2011 which documented the flooding were selected. A list of Memphis road closures on 12 May 2011 ($n = 37$) was collected from an on-line news source. Some of the data contained geolocation information, while others were geolocated using the Google API.

2.2 Remote sensing data

Full-resolution GeoTIFF Multispectral Landsat ETM+ images for 2 January and 10 May 2011 are used. The data were downloaded from the USGS Hazards Data Distribution System (HDDS). Landsat data are comprised of seven spectral bands: optical (0.45–0.52, 0.52–0.60, 0.63–0.69 μm), near-IR (0.77–0.90 μm), mid-IR (1.55–1.75, 2.09–2.35 μm) and thermal-IR (10.40–12.50 μm) with a spatial resolution of 30 m. The images were georeferenced to UTM coordinates in ArcGIS and an area encompassing Memphis and its greater metropolitan area was selected at a scale of 1:145 000.

2.3 Digital elevation model data

A USGS Seamless Data Warehouse DEM with a 30 m resolution was used in this study. The DEM is georeferenced to UTM coordinates in ArcGIS and exported at the same 1:145 000 scale as the Landsat data (Fig. 1).

2.4 Meteorological data

Meteorological data relative to maximum daily precipitation rate and total daily precipitation were obtained from the

¹www.r-project.org

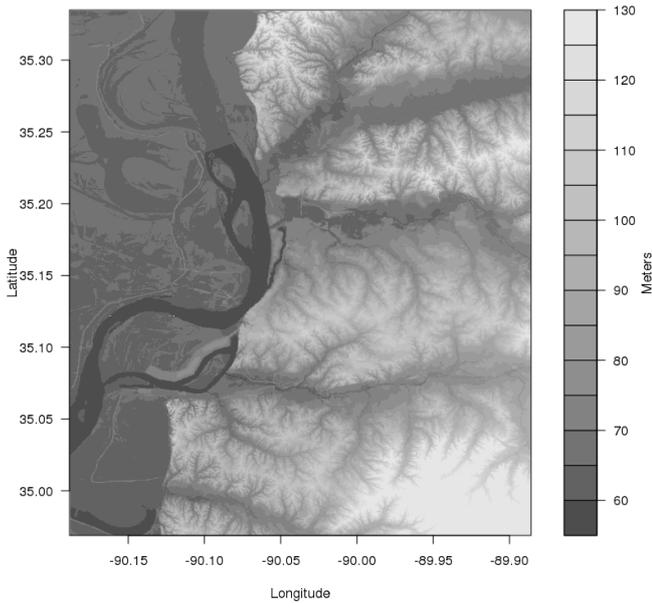


Fig. 1. Digital Elevation Model for the region of study.

NCEP CPC Morphed Precipitation Model (Joyce et al., 2004) and Weather Underground² (WU), respectively. Figure 2 shows the NCEP daily precipitation rate (bars) and the WU accumulated precipitation (solid line) for the period ranging from 1 April to 31 May 2011. The acquisition time for the May Landsat data is shown, and it occurs after the period of intense rainfall during the end of April. These meteorological data are used to identify appropriate dates for the remote sensing data. It is desirable that a scene is selected after a period of intense rainfall in order to identify the maximum flood extent.

2.5 River gage data

River gage data for the Mississippi River in Memphis were collected from the US Army Corps of Engineers RiverGages³ website. The data used for this study were collected from gage MS126 located at longitude: 90.07667000° W, latitude: 35.12306000° N. Data were selected in elevation (m) format so they could effectively be used in conjunction with the DEM.

Figure 3 shows the height information for MS126 for the entire year 2011. The acquisition time for the January and May Landsat data are indicated, and they correspond, respectively, to the almost minimum and maximum water heights for the entire year. The river gage height information is paired with the DEM to derive the approximate flood extent.

²www.weatherunderground.com

³www.rivergages.com

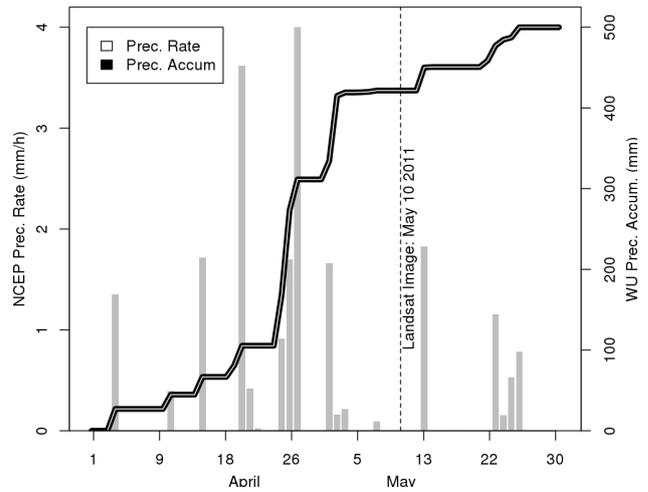


Fig. 2. Maximum daily precipitation rate and accumulated precipitation for the period ranging from 1 April to 31 May 2011.

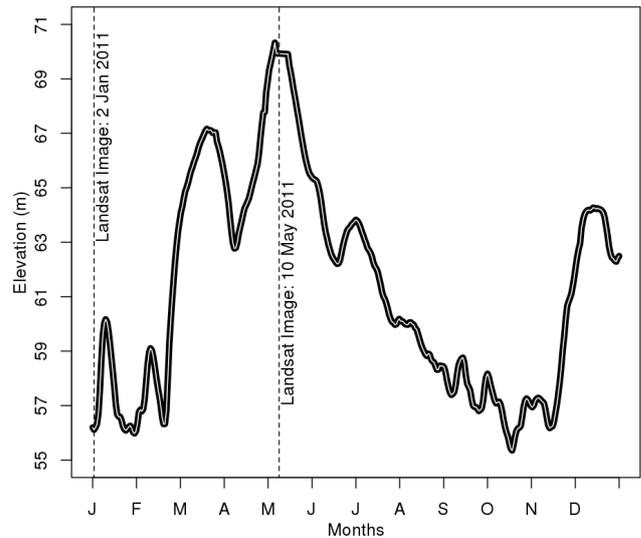


Fig. 3. Year 2011 water height profile for Mississippi River at Memphis, TN.

3 Methodology

3.1 Overview

The proposed methodology is based on the data fusion of different layers generated from different data sources. Figure 4 illustrates this integration of multiple layers, which may have varying resolutions or sparse data. The output is shown in the bottom most layer, where a flood hazard map is generated. The input consists of different layers generated using remote sensing data, DEM, ground information, etc., as normally discussed in the literature. The novelty of the proposed methodology, is the introduction of volunteered data as an additional layer, and their use in refining the

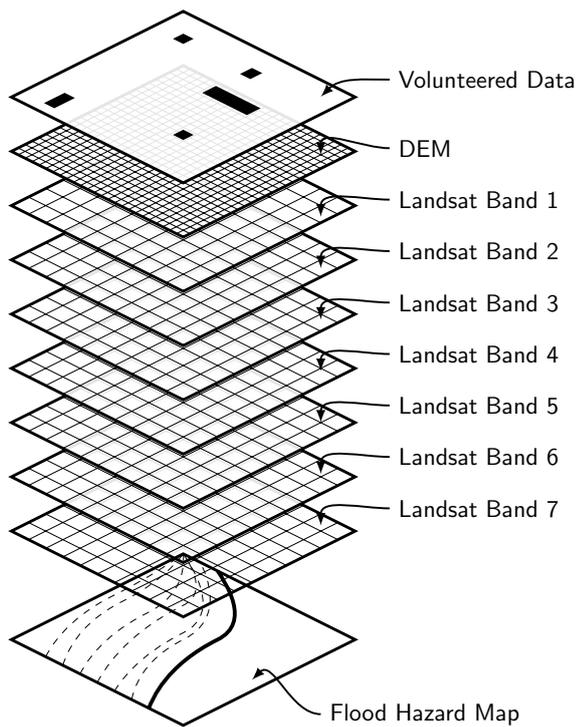


Fig. 4. The proposed methodology fuses several layers to generate a flood hazard map.

hazard map. Therefore, although in this paper we used specific remote sensing and DEM data, the methodology is not limited by these data types and can easily be extended to integrate additional or different sources. It is crucial for obtaining the most accurate measurements that the correct classification methodology is applied to each data type when creating the flood extent layer(s).

The ground data used is not a verified “ground truth”, but can be utilized as reliable information to assess the presence or absence of water in specific locations. It compensates for the potential misclassification of remote sensing data due to resolution, satellite orbit limitation, cloud cover, or data acquisition problems. Furthermore, the volume of the data alone as a function of time can be an indication of the geospatial rate of progression of the event, and can help prioritize response to specific areas.

The methodology consists of a three step process:

1. Identification of flood extent.
2. Generation of flood hazard maps.
3. Ground data integration.

3.2 Identification of flood extent

Different methodologies can be used to identify the extent of water over the geographical region of interest. The goal of this step is to generate one or more maps using the input

layers which identify regions where water is detected. The task is method independent, and it can use any method that is best suited for a particular combination of data and location.

In this article, two different methods are employed to identify flood extent. The first involves the use of remote sensing data and machine learning classification, and the second involves the use of a DEM and river gage data.

3.3 Generation of flood hazard maps

After one or more flood extent maps are generated, a flood hazard map is created by computing the probability for each pixel to be flooded. This probability map is generated by applying a kernel density smoothing operation over the 2-D data, and then by normalizing the result. Let $(w_1x_1, w_1x_2, \dots, w_1x_n)$ be weighted samples drawn from a distribution with an unknown density f , the goal is to estimate the shape of this function. The general kernel density estimator is

$$f(x) = \frac{1}{nh} \sum_{i=1}^n K\left(\frac{w(x-x_i)}{h}\right), \quad (1)$$

where K is the kernel function, h is the bandwidth, and w is a user selected weighting scalar. The weight w describes the importance of a particular observation, or the confidence associated with the flood extent map. In this application, using a weighted kernel function is paramount because ground observations cannot be considered “ground truth” proper, since volunteered geographical data carry intrinsic uncertainties due to their generally non-scientific nature. Therefore, using different values of w properly assigns levels of confidence to the various observations.

The identification of weight w is problem specific and domain dependent, but most importantly, it is dependent on data quality. A weight is used to include the concept of “significance” of the data in the algorithm and the analysis. It is assumed when working with such heterogeneous data, the information might vary significantly, and therefore decisions should be made, when possible, using better data.

Quantitative measures to define the weight can be established. For example, when using satellite data the pixels along the center of the swath or those that are cloud free are preferred in the analysis. Most satellite products have a quality index associated with each pixel that can be used to set the appropriate weight.

For volunteered data, the weight may vary depending on the characteristics of the source. For example, the volume of the data can be used to assign higher weight to data with dense spatial coverage and numerous observations. Higher weight can also be dependent on the source itself. For example, observations published by local news are assumed to have been validated more than points volunteered anonymously. Finally, there is also a subjective component that can be taken into account, assigning different weight to specific users or regions.

The output of the kernel smoothing is a map with contour lines illustrating the probability that specific regions are flooded. The specific methodology used for Kernel smoothing and its R implementation used in this study are described by Wand and Jones (1995).

3.4 Ground data integration

The last step consists of modifying the flood hazard map by the integration of the ground data. The nature of this data is different from the data used to generate the previous flood extent maps. It is usually comprised of sparse point data, identifying the presence or absence of flooding in a specific region.

For this study, weight values are assigned experimentally. They are first equally assigned and then adjusted based on characteristics and confidence in the data source. By assuming the machine learning tree induction and the DEM/river gage approach are equally adept at classifying water, their weight values are kept constant while values of the volunteered ground data are adjusted. The flooded roads documented by local news sources are assumed more reliable than the sparse, point data of the videos and the pictures. Based on this assumption, the weights were set to 3 for the news data, and to 2 for the pictures and videos. Equal values of 1 were assigned for both the DEM and Landsat data.

4 Results

4.1 Flood classification using DEM and river gage data

A DEM and river gage data are used to classify water pixels for 2 January and 10 May 2011. Pixels in the DEM with a height less than or equal to the river gage height are set as water pixels. Specifically, heights of 56 m and 70 m are used for January and May, respectively. Figure 5a, c show the areas identified as water for January and May dates, imposed over the DEM. The scale information is the same as in Fig. 1.

4.2 Flood classification using machine learning tree induction

Water pixels are identified in both the January and May Landsat images by using a machine learning tree induction classifier. Ripley (2008) describes the general rule induction methodology and its implementation in the R statistical package used in this study.

For the machine learning classification, 4 control areas of roughly the same size are identified, 2 over the Mississippi river as examples of water pixels and 2 over different regions with no water pixels as counter-examples. Landsat multispectral data relative to these regions are used as training events by the decision tree classifier. The learned tree is then used to classify the remaining water pixels in the scene. This process is repeated for both the January and May scene, and is illustrated in Fig. 5b, d.

About 1% of the total number of pixels are used as training pixels (events), and the remaining 99% are classified according to the induction tree generated.

4.3 Flood hazard maps and ground data integration

The methodology described in Sect. 3.3 is employed to generate flood hazard maps using both the DEM and Landsat pixel classifications. The goal is to assign each pixel a probability of being part of the flooded area. Figure 6a, c show the probability contour lines for January and May, respectively.

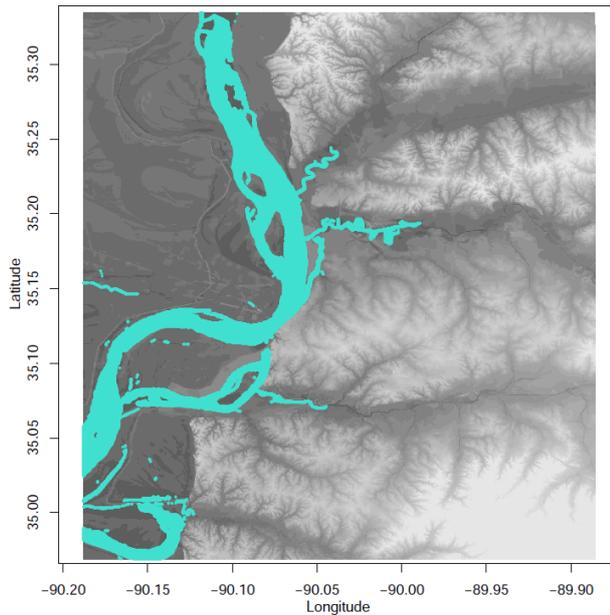
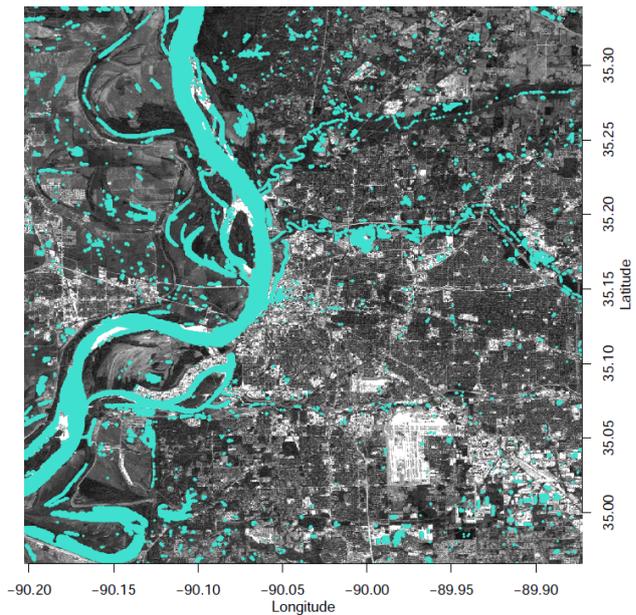
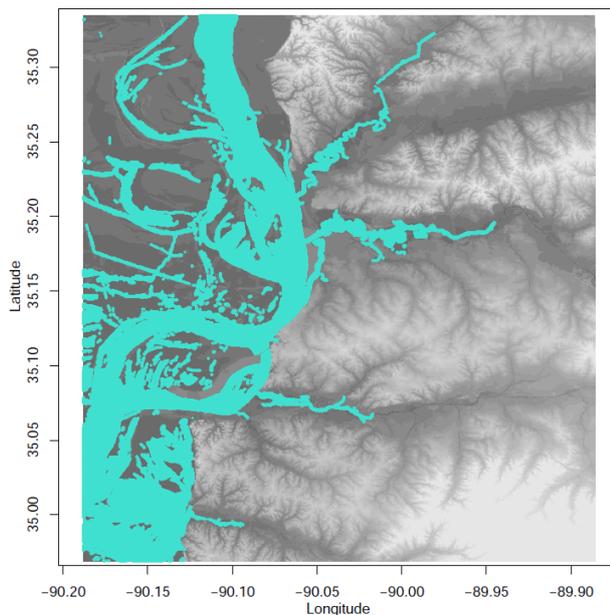
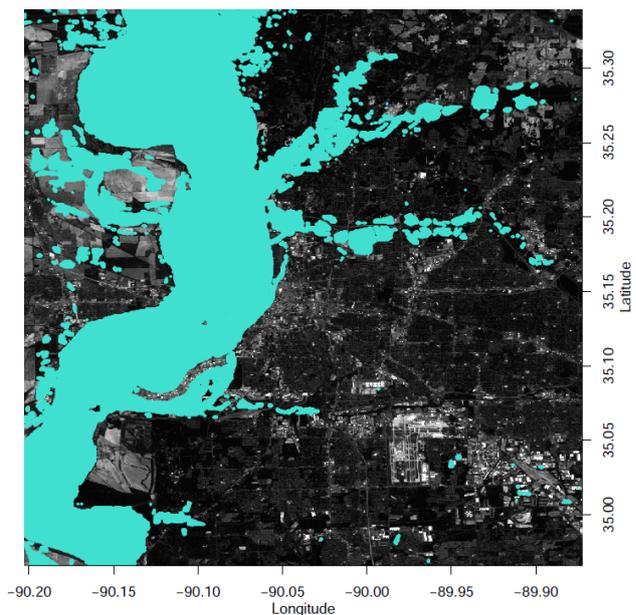
Additional data with ground information is then used to refine the January and May flood hazard maps (Fig. 6b, d). The images show both the location and type of the ground information (video, photos, news), and the resulting probabilities when these data are taken into account. The data are imposed on both the January and May hazard maps, although all ground data are collected from the May flood. The image generated for the May flood, (Fig. 6d), shows modifications to the flood hazard map after the incorporation of the ground data. The ground data are also incorporated with the January (non-flood) image (Fig. 6b) to illustrate how a preliminary hazard map could be generated if current satellite data are not available. In both instances, the addition of supplemental information in the form of volunteered ground data alters the flood map by expanding the area of possible inundation and by adjusting pixel values.

The pixel classifications are summarized in Table 1 and by a histogram in Fig. 7. As expected, when generally comparing flood versus non-flood scenes as in Fig. 6a, c, more pixels have a higher probability (60–100%) of being flooded and less pixels have a lower probability (0–40%) of being flooded in the May (flooded) image as compared to the January (non-flooded) image.

When the ground data are incorporated into the hazard maps (Fig. 6b, d), a spatial analysis shows noteworthy changes. The incorporation of ground data yields enhancements to the May flood hazard map (Fig. 6d) which are evident by the progression of contour lines and reclassifications of pixels. Higher value contour lines, indicating a greater probability of a region being flooded, progress toward the northeast, where the majority of ground information are located. Examining the differences between the two May scenes in Table 1, the percentage of pixels classified as having a low probability (0–20%) of being flooded as well as the pixels classified as having a high probability (80–100%) of flooding decreases or increases 6 percentage points, respectively, after the incorporation of ground data. These changes illustrate that although both the DEM/river gage and Landsat classification techniques can be highly accurate in identifying flooded areas, the addition of real-time on the ground data, verifying the presence of water in a specific area, can augment an inundation map.

Table 1. Number of pixels classified as water.

	P(w)	Jan	Jan + Ground	May	May + Ground
0–20 %	126 567 (51 %)	108 932 (44 %)	114 104 (46 %)	99 368 (40 %)	
20–40 %	53 718 (21 %)	57 475 (23 %)	40 190 (16 %)	30 373 (12 %)	
40–60 %	21 136 (08 %)	27 601 (11 %)	21 235 (08 %)	26 997 (11 %)	
60–80 %	12 660 (05 %)	18 969 (08 %)	23 260 (09 %)	29 368 (12 %)	
80–100 %	35 919 (14 %)	37 023 (15 %)	51 211 (20 %)	63 894 (26 %)	

**(a)** Water Classification using DEM for Jan 2011**(b)** Water Classification using Landsat for Jan 2011**(c)** Water Classification using DEM for May 2011**(d)** Water Classification using Landsat for May 2011**Fig. 5.** Water pixel classification using the DEM (**a** and **c**) and Landsat (**b** and **d**) and for January (top) and May (bottom) data. The background (**b** and **d**) is from Landsat band 3.

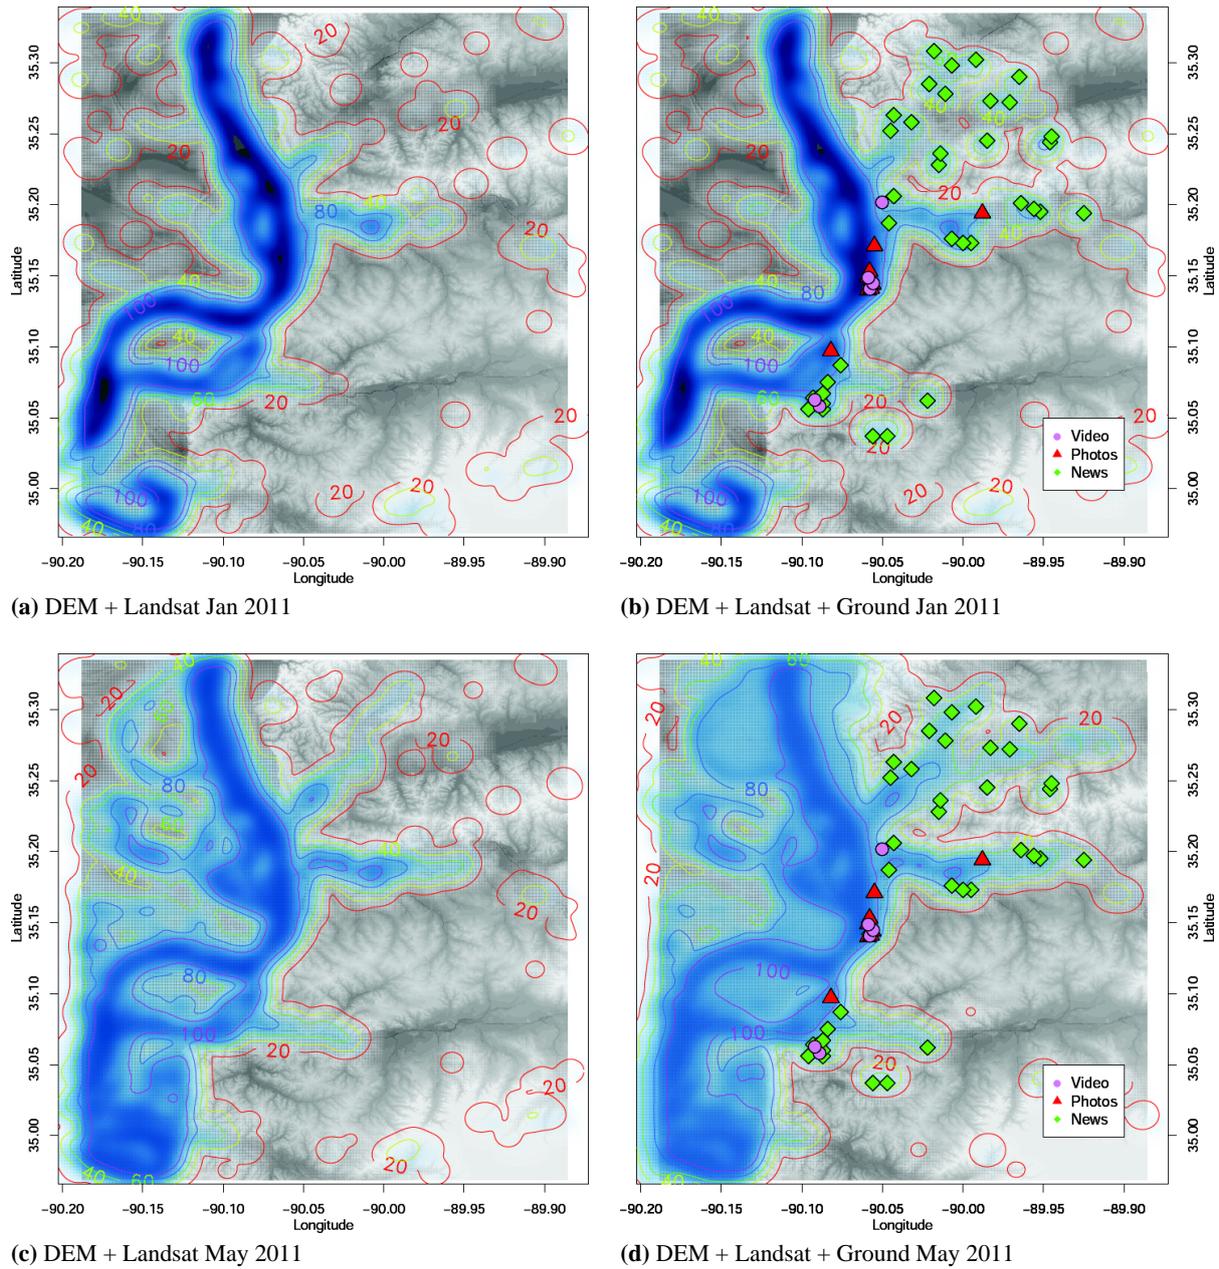


Fig. 6. Flood hazard map indicating the probability of flood in percentage using DEM, Landsat, and ground data for January 2011 (a and b) and for May 2011 (c and d).

Applying the layer of ground data to the January Hazard Map (Fig. 6b) illustrates how a small amount of real-time volunteered ground data could be integrated with an historical image to identify possible flooded regions. The number of non-water pixels (0–20%) is reduced by 7 percentage points and reclassified to higher probability classes (Table 1).

Figure 6b,d show while volunteered ground data does modify the flood hazard maps, the amount of modification is limited by the spatial distribution of the ground data. The evolution of the contour lines, or areas of change, in both images are restricted to regions where the volunteered data are located. This illustrates while the incorporation of volunteered ground data does affect a change in both flood hazard maps, the areas of change are controlled by the quantity and distribution of the volunteered data.

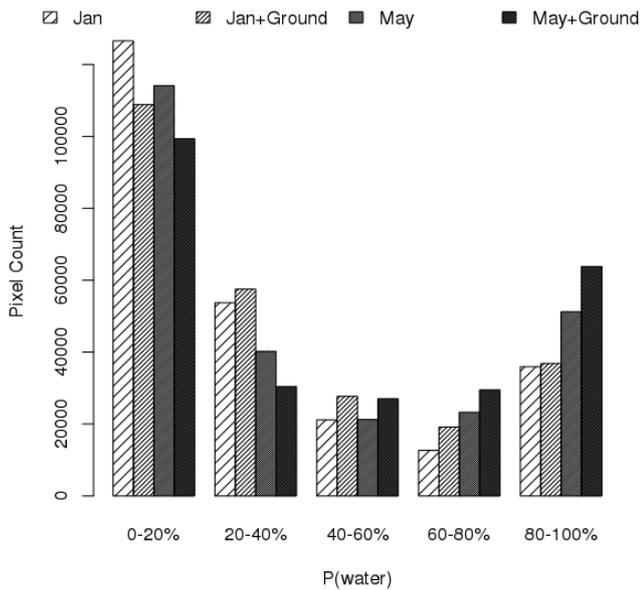


Fig. 7. Number of pixels classified as water.

5 Conclusions

This paper presents a new methodology fusing volunteered geographical data with remote sensing and DEM data for the creation of hazard maps. It is shown that even a small amount of ground data points can change the flood assessment when compared to satellite and DEM data alone. The spatial distribution of the volunteered data limits the areas where change is detected.

The methodology is particularly useful when satellite data is limited or of poor quality. Additionally, the ground data can add a time component which can help determine the change occurring between remote sensing observations.

Acknowledgements. Work performed under this project has been partially supported by the US Department of Transportation award 202717 (RITARS-12-H-GMU, CFDA).

The publication of this article was funded in part by the George Mason University Libraries Open Access Publishing Fund.

Edited by: N. Kerle

Reviewed by: three anonymous referees

References

- Acar, A. and Muraki, Y.: Twitter for crisis communication: lessons learned from Japan's tsunami disaster, *International Journal of Web Based Communities*, 7, 392–402, 2011.
- Apel, H., Thielen, A., Merz, B., and Blöschl, G.: A probabilistic modelling system for assessing flood risks, *Nat. Hazards*, 38, 79–100, 2006.

De Longueville, B., Smith, R., and Luraschi, G.: “OMG, from here, I can see the flames!”: a use case of mining location based social networks to acquire spatio-temporal data on forest fires, in: *Proceedings of the 2009 International Workshop on Location Based Social Networks*, 73–80, 2009.

Earle, P., Bowden, D., and Guy, M.: Twitter earthquake detection: earthquake monitoring in a social world, *Ann. Geophys.*, 54, 708–715, 2012, <http://www.ann-geophys.net/54/708/2012/>.

Flanagin, A. and Metzger, M.: The credibility of volunteered geographic information, *GeoJournal*, 72, 137–148, 2008.

Frazier, P. and Page, K.: Water body detection and delineation with Landsat TM data, *Photogramm. Eng. Rem. S.*, 66, 1461–1467, 2000.

Goodchild, M.: Citizens as sensors: the world of volunteered geography, *GeoJournal*, 69, 211–221, 2007.

Heverin, T. and Zach, L.: Microblogging for crisis communication: examination of Twitter use in response to a 2009 violent crisis in Seattle-Tacoma, Washington area, in: *Proceedings of the Seventh International Information Systems for Crisis Response and Management Conference (ISCRAM)*, Seattle, Washington, 2010.

Hyvärinen, O. and Saltikoff, E.: Social media as a source of meteorological observations, *Mon. Weather Rev.*, 138, 3175–3184, 2010.

Joyce, R., Janowiak, J., Arkin, P., and Xie, P.: CMORPH: A method that produces global precipitation estimates from passive microwave and infrared data at high spatial and temporal resolution, *J. Hydrometeorol.*, 5, 487–503, 2004.

Laura, L., Melack, J., and David, S.: Radar detection of flooding beneath the forest canopy – A review, *Int. J. Remote Sens.*, 11, 1313–1325, 1990.

Liu, S., Palen, L., Sutton, J., Hughes, A., and Vieweg, S.: In search of the bigger picture: the emergent role of on-line photo-sharing in times of disaster, in: *Proceedings of the Information Systems for Crisis Response and Management Conference (ISCRAM)*, 2008.

McDougall, K.: Using volunteered information to map the Queensland floods, in: *Proceedings of the Surveying and Spatial Sciences Biennial Conference*, 2011.

Poser, K. and Dransch, D.: Volunteered geographic information for disaster management with application to rapid flood damage estimation, *Geomatica*, 64, 89–98, 2010.

Pultar, E., Raubal, M., Cova, T., and Goodchild, M.: DDynamic GIS case studies: wildfire evacuation and volunteered geographic information, *Transactions in GIS*, 13, 85–104, 2009.

Richter, B., Baumgartner, J., Braun, D., and Powell, J.: A spatial assessment of hydrologic alteration within a river network, *Regul. River.*, 14, 329–340, 1998.

Ripley, B.: *Pattern recognition and neural networks*, Cambridge University Press, 2008.

Sanyal, J. and Lu, X.: Application of remote sensing in flood management with special reference to monsoon Asia: a review, *Nat. Hazards*, 33, 283–301, 2004.

Schlieder, C. and Yanenko, O.: Spatio-temporal proximity and social distance: a confirmation framework for social reporting, in: *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on Location Based Social Networks*, ACM, 60–67, 2010.

- Smith, L.: Satellite remote sensing of river inundation area, stage, and discharge: A review, *Hydrol. Process.*, 11, 1427–1439, 1997.
- Tapia, A., Bajpai, K., Jansen, B., Yen, J., and Giles, L.: Seeking the trustworthy Tweet: can microblogged data fit the information needs of disaster response and humanitarian relief organizations, in: *Proceedings of the 8th International Conference on Information Systems for Crisis Response and Management (ISCRAM)*, 10 pp., 2011.
- Townsend, P. and Walsh, S.: Modeling floodplain inundation using an integrated GIS with radar and optical remote sensing, *Geomorphology*, 21, 295–312, 1998.
- Tyshchuk, Y., Hui, C., Grabowski, M., and Wallace, W.: Social media and warning response impacts in extreme events: results from a naturally occurring experiment, in: *Proceedings of the 45th Annual Hawaii International Conference on System Sciences (HICSS)*, IEEE, 818–827, 2012.
- Verma, S., Vieweg, S., Corvey, W., Palen, L., Martin, J., Palmer, M., Schram, A., and Anderson, K.: Natural language processing to the rescue? Extracting “situational awareness” Tweets during mass emergency, in: *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work Companion*, 227–230, 2011.
- Vieweg, S., Hughes, A., Starbird, K., and Palen, L.: Microblogging during two natural hazards events: what twitter may contribute to situational awareness, in: *Proceedings of the 28th international Conference on Human Factors in Computing Systems, ACM SIGCHI*, 1079–1088, 2010.
- Wand, M. and Jones, M.: *Kernel smoothing*, Vol. 60, Chapman & Hall/CRC, 1995.
- Wang, Y., Colby, J., and Mulcahy, K.: An efficient method for mapping flood extent in a coastal floodplain using Landsat TM and DEM data, *Int. J. Remote Sens.*, 23, 3681–3696, 2002.
- Zhang, H., Korayem, M., Crandall, D., and LeBuhn, G.: Mining photo-sharing websites to study ecological phenomena, in: *Proceedings of the 21st international conference on World Wide Web*, 749–758, 2012.