Mining and Classifying Image Posts on Social Media to Analyse Fires

Shannon Daly
RMIT University, Melbourne, Australia
s3285539@student.rmit.edu.au

James A. Thom
RMIT University, Melbourne, Australia
james.thom@rmit.edu.au

ABSTRACT

We propose a methodology to study the occurrence of fires through image posts on Flickr; crowd-sourcing information from a noisy social media dataset can estimate the presence of fires. We collect several years worth of photos and associated metadata using fire-related search terms. We use an image classification model to detect geotagged photos that are further analysed to determine if a fire event did occur at a particular time and place. Furthermore, a case study investigates image features and spatio-temporal elements in the metadata, as well as location information contained in camera EXIF data.

Keywords

Flickr, image analytics, geotags, geocoding.

INTRODUCTION

Photos submitted by users on social networks have led to the formation of massive online photo collections. The image-hosting site Flickr[1] has become considerably prominent; as of 2015 it contains over 10 billion photos[2] and is estimated that users upload over 1.8 million photos per day[3]. Not only does Flickr contain a vast collection of photos, it also holds a wealth of metadata associated with each photo including temporal and spatial information. Mining metadata on social media has even produced effective real-time monitoring systems that are able to detect potential natural disasters (Power, Robinson and Ratcliffe, 2013; Robinson, Power and Cameron, 2013). Large collections of photos and metadata, however, become problematic for analysis. Data is often noisy and photos may contain poor annotation (or no annotation at all). Thus, textual attributes associated with an image such as titles, descriptions and tags may not allow for accurate filtering.

In a crisis, there is an increased need to communicate with others. Users often seek information on the current status of the disaster and share information themselves. With photo-sharing social networks, there is an additional benefit of being able to view and post useful imagery to enhance situational awareness, allowing users, emergency services and local authorities to assess the scope and damage of an affected region. Flickr users often capture and share photos when a natural disaster occurs, providing a visual snapshot of the events that took place and how they developed over time. Some of these photos contain location information.

In this research, we aim to investigate visual characteristics and spatio-temporal distributions of major fires that have occurred based on a noisy Flickr dataset. Although many photo-sharing sites exist, Flickr was chosen as it

[1] http://www.flickr.com

[2] http://blog.flickr.net/en/2015/05/07/flickr-unified-search/

[3] https://www.flickr.com/photos/franckmichel/6855169886

contains a large amount of high-quality images. Flickr also publishes EXIF header data associated with each image uploaded, enabling further temporal and geospatial analysis of the data.

Our first task identifies photos containing instances of fire. While methods exist to detect fire and smoke in real-time video surveillance, less attention has been applied to identifying these concepts in a dataset comprised of still images. Although this is a simple task for humans to recognise and identify the presence of fire and smoke in an image, it is much more challenging when attempting to automate this process digitally. For example, images do not allow for video processing techniques such as motion detection and change detection.

Our second task measures the presence of a fire on social media based on Flickr photos and metadata. We select a region where a fire has occurred and then specify an appropriate time period to restrict our analysis. We then quantify and measure the amount of geotagged photos posted from that region and compare with ground truth data of the fire, that allows for an evaluation of the location of geotags against the actual fire impact area.

As a final task, we investigate if location information available in EXIF headers on Flickr is an accurate source. So far, little to no attention has been paid to this information. This may be in part due to the low amount of photos containing location information in EXIF headers. However, we discover a sufficient portion of photos from our Flickr dataset contain them, enabling us to perform an evaluation via a case study.

Our motivation for this research is two-fold. Major fires cause a large impact in affected regions and, while there is work that has produced effective solutions to monitor fires on social media using text-based content, less attention has been paid to the visual content. Crowd-sourcing information from social media and gathering relevant images allows for a visual assessment of a fire that has occurred. Secondly, while images of a fire may be attainable through artificial satellites orbiting Earth, they may not provide an accurate visual assessment of conditions on the ground. One particular challenge in aerial views of a fire is that smoke and clouds may present visual ambiguity. There appears to be an untapped resource of rich imagery on social media that may assist in post-hoc research of a major fire event. Images posted on social media, such as Flickr, may fill gaps where information is not available.

BACKGROUND AND RELATED WORK

We first overview the specific image classification techniques we use in our research, then present a survey of literature related to geo-spatial analysis.

Object Detection and Image Classification Techniques

Automated image analysis and object detection has long been studied before the advent of social media. A breakthrough in this field was first by Lowe (1999). The algorithm, known as Scale-Invariant Feature Transform (SIFT), is widely used to recognise objects by identifying interesting sections of an image to extract features. Based on these located features, it is then able to find the same object within another given image. ColorSIFT, or CSIFT, (Abdel-Hakim and Farag, 2006), was introduced as a more robust model; rather than consider images in a greyscale spectrum, as in the original SIFT implementation, the ColorSIFT model builds the descriptors in a color invariant space.

While the above models are effective, there is a caveat whereby training a system with these models may be time-consuming. An alternative approach, Speeded Up Robust Features (SURF) was presented by Bay, Ess, Tuytelaars and Van Gool (2008) as a way of reducing computation time by using existing descriptor models, however simplifying their characteristics to a minimum. Akin to SIFT descriptor models, the SURF model was also extended to consider color (Fan, Men, Chen and Yang, 2009). A different approach, GIST (Oliva and Torralba, 2001), rather than process individual regions of an image, takes a holistic view. GIST effectively allows for the recognition of a scene in an image by accounting for dimensions such as naturalness, openness, roughness, expansion and ruggedness to generate a multidimensional space that can be used to determine the semantic category.

As the proliferation of images on social media has risen, algorithms such as these have been invaluable to the research community in order to analyse and classify images. However, not all approaches have considered visual features alone. The analysis of images typically falls into two categories: image analysis based on visual features, which assesses the images at a low level and image analysis based upon non-visual features; in the context of social media, this includes metadata such as geotags, descriptions, tags and other attributes.

Li and Fei-Fei (2007) proposed the use of the SIFT algorithm in event categorisation, and using SIFT descriptors to label images achieved a system of classifying images with 73.4% accuracy. Cao, Luo, Kautz and Huang (2008) aimed to annotate photos by exploring the GPS and timestamp of photos, and they used this metadata to effectively label a scene. An image analysis system to recognise landmarks was developed by Li, Crandall and Huttenlocher (2009). Their dataset consisted of over 30 million geotagged photos that were posted on Flickr, making it one of the largest datasets of its kind. Their approach consisted of training their data using a multiclass SVM and subsequently analysing images using the SIFT algorithm. Results showed that classification accuracy was comparable to that of manual classification performed by a human. The paper clarified that automated image classification is quite efficient on such a large dataset and can reduce the need for manual processes greatly. More recently, Amato, Falchi and Bollitieri (2010) performed an evaluation of visual features in the ability to recognise landmarks as well. Their dataset however was significantly smaller, involving only 1,227 Flickr photos of landmarks located in Pisa, Italy. Results showed that local features were more accurate than global features and the SIFT algorithm was more accurate than SURF and ColorSIFT counterparts.

Chen and Roy (2009) noted that capturing the visual features of photos to classify them was difficult, and instead classified Flickr photos by exploiting the use of *tags* supplied by users for each photo. Their approach identified tags that were related to events and then subsequently identified photos corresponding to each cluster of events. Firan, Georgescu, Nejdl and Paiu (2010) extended upon this approach by analysing images using tags, titles and photo descriptions to classify them more accurately. Their approach to cluster Flickr images had high success by grouping events listed on Wikipedia with images on Flickr groups. McAuley and Leskovec (2012) similarly took advantage of metadata such as comments, geotags, friends and associated image groups to classify images. The focus of classifying images using metadata was important in these studies as the lack of geotags is a problem; in some datasets it has been found only 1% of images have been geotagged (Middleton, Middleton and Modafferi, 2014). However, as the proliferation of GPS-enabled devices has risen steadily over recent years the amount of geotagged posts has risen and thus warrants further investigation. Image analysis of photos posted on Twitter was performed by Chen, Lu, Kan and Cui (2013) by extracting SIFT descriptors from the images and subsequently clustering them to form visual words. Text features were additionally taken into account by way of a word segmenter. Their dataset consisted of over 57 million tweets, with 45.1% of them being image tweets. The accuracy of image classification using text, image and social context features was reported as a Macro-F1 score of 70.5%.

Fire and Smoke Detection

The ability to detect a hazard such as fire and smoke is not new. Hazards such as these can cause large devastation and financial loss and hence the ability to detect them as early as possible are being sought after. The use of social media by millions of users has seen an influx of valuable information provided during a natural disaster. Photos of disasters such as fires are often taken and shared on online photo-sharing social networks. News and media often include photo and video taken by users in their broadcasts to show ongoing conditions on the ground. However, this media is often hard to find. Social networks contain noisy datasets; photos often include poor annotations. Detection of fire and smoke in image and video has provided a means to overcome this problem. Hadjisophocleous, Ouyang, Ding and Liu (2010) studied the performance of the AlarmEye VID smoke and flame detection system. Their tests involved creating prepared fire sources in a large atrium facility and evaluating the effectiveness of the detection system. Even under challenging conditions, fires and smoke was detected well. Classifying fire pixels was proposed by Çelik, Özkaramanli and Demirel (2007) using fuzzy logic. Their method is able to effectively separate fire and non-fire pixels by using the *YCbCr* color space to separate luminance from chrominance. Their approach achieves a correct detection rate of 99%, with a false alarm rate of 9.5%. However, their model can only be used as a pre-processing method in fire detection systems that analyse video sequences.

Geo-spatial Analysis on Social Media

Geo-referenced photos on social media may contain location information that originates from two sources; geographic attributes contained in external metadata tags and those found in EXIF header data. In this research, *geotags* refers to the geographic location where an image was uploaded as contained in external metadata tags, whereas the term *geocoded* refers to an image with geographic attributes[4] contained in the EXIF header data.

Exploring geotags associated with images posted on social networks has identified applications in image classification. Huang, Cheng, Hsie, Hsu and Chang (2011) used a model to recognise famous landmarks by crowdsourcing geotagged Flickr photos. They achieved this by selecting a region and using geotagged photos within that region to extract landmark names from Flickr tags. However, problems related to geotags were also evident; low F1 scores were recorded due to people taking photographs at a large distance away from the landmark.

Zhang, Korayem and Crandall (2012) performed a study on mining social media to predict the presence of snow based on geotagged Flickr photos. In their paper, a method they use to predict snow in a given geospatial area is by analysing the amount of geotagged photos containing snow in a defined region. Using this technique allows them to effectively identify regions where snow and other ecological phenomena have occurred.

Middleton et al. (2014) created real-time crisis maps – maps depicting the geospatial coverage of a natural disaster – by analysing tweets on Twitter. They performed an evaluation of geotagged tweets against ground truth data to then evaluate their crisis maps. To do this, they plotted locations on a map and calculated precision, recall and F1 measures by segmenting their map into a grid of cells and identifying true positives, true negatives, false positives and false negatives. True positives were reported as cells that contain both geotags and ground truth storm activity. They demonstrated a high precision (90+%) and show that it is possible to use location information from Twitter to achieve this.

METHODOLOGY

Our methodology comprises six key tasks performed to analyse fires on photo-sharing sites:

- Obtaining our main dataset of 114,098 images for analysis of fires
- Obtaining and manually classifying a separate dataset of training and testing images
- Building and testing an image classifier for ‘fire’
- Classifying images in our dataset
- Identifying major fires based on ‘fire’ images
- Performing a case study of the California Rim Fire

Main Dataset

A sample of 114,098 publicly-accessible photos were collected using the Flickr API by performing repeated searches using the *flickr.photos.search* method, using the terms “bushfire” and “wildfire”. Our search limited the time period to photos between January 1, 2009 to October 1, 2015 by specifying a *min_upload_date* timestamp in our search parameters (see Figure 1). This restriction was performed to eliminate conflicting occurrences related to training and testing images. Each image and its associated metadata was stored, with metadata being stored and organised in a MySQL database. Photos with and without geo-spatial data were collected, and EXIF data is collected using the *flickr.photos.GetEXIF* API method. Overall, the total size of images collected from Flickr was almost 13GB, and the size of our database was almost 2.5GB.

[4] http://www.exif.org/Exif2-2.PDF

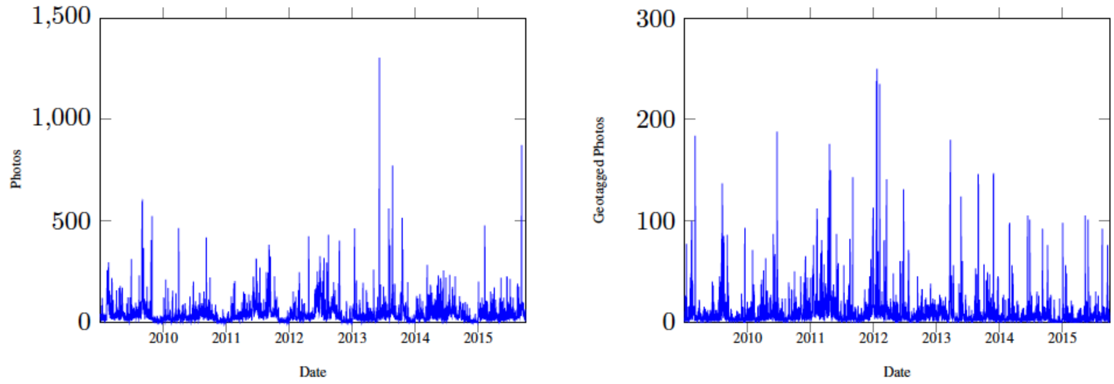

Figure 1. Daily count of all images (left), and only geotagged images (right) in our collection, from January 1, 2009 to October 1, 2015.

Image Classification Model

In the following, we discuss our approach to build our image classifiers. We begin with our method of collecting training and testing images. Then describe the process we use to build our image classifiers based on our training and testing images obtained. Finally, we present the results of our image classification performance.

Image dataset

In order to build the image classification system, an array of training and testing images were needed that related to fire activity. Flickr photos were also used as a source for training and testing:

- training images were collected from February 2004 (Flickr website launch) to December 2007,
- test images were collected from January 2008 to December 2008.

Photos were downloaded at a medium resolution (no Flickr URL suffix[5] set). Described below are the methods and parameters specified in gathering training images related to each subject.

Images of fire were collected using the search term “fire” and were manually judged, 974 images were used for training our ‘fire’ classifier, consisting of 674 positive and 300 negative samples. Images were manually cropped when necessary to display only the subject. 505 images were collected for testing the positive detection rate and 249 images were gathered to test the false positive rate.

Figure 2. Samples of positive training images collected from Flickr that were used to build our binary image classifier for ‘fire’.

[5] https://www.flickr.com/services/api/misc.urls.html

Building Image Classifiers

In order to cluster our photos into groups for further investigation, we analyse visual features of each photo. For image classification, we apply the Bag of Features model described by Csurka, Dance, Fan, Willamowski and Bray (2004).

Our first step in building our classifiers is obtaining keypoints by using a feature detector. Using this, we extract descriptors from each keypoint. Vector quantisation is then performed by clustering our set of descriptors using the *k-means* algorithm to calculate the Euclidean distance between each cluster center and each feature to identify the closest point to each cluster. Our visual vocabulary is then generated using these clusters. Once our visual vocabulary is obtained, keypoints are detected for each image and used to extract associated descriptors. Each of these descriptors is matched to its nearest neighbor in our visual vocabulary by implementing the Fast Library for Approximate Nearest Neighbors (FLANN) (Muja and Lowe, 2009). A histogram of k features is then generated (using k = 1000 performed well). We then learn a Support Vector Machine (SVM) binary image classifier. We perform 10-fold cross-validation on our binary image classifier for optimal parameter selection. We use a Gaussian RBF kernel given by

$$K(x_i, x_j) = e^{-\gamma\|x_i - x_j\|^2}, \ \gamma > 0$$

where:

- x_i and x_j are our feature vectors in the feature space,
- γ influences the reach of a single training sample, and
- $\|x_i - x_j\|^2$ is the Squared Euclidean Distance of our feature vectors.

The regularisation parameter C is also selected during cross-validation, where C is the trade-off value between misclassification of training samples against the model complexity. To calculate each classifier's performance, we run each of them through our positive and negative test images.

Performance

For classifying fire, we compare four different approaches (utilising the Bag of Features model and the same training and testing data) to train the SVM: SIFT, SURF, ColorSURF and ColorSIFT. We found ColorSIFT performed best, as shown in Table 1, giving a 91% recall and 93% precision.

	Fire	Not Fire
Classifier predicts Fire	460	33
Classifier predicts Not Fire	45	216

Table 1: Confusion matrix for performance of classifier with ColorSIFT feature model on the test data set, showing number of true positives (TP=460), false positives (FP=33), false negatives (FN=45) and true negatives (TN=216).

Figure 3 shows examples of a true positive, a false positive, a false negative, and a true negative from the test dataset using the classifier trained with ColorSIFT feature model.

True Positive

False Positive

False Negative

True Negative

Figure 3. Samples of TP, FP, FN, TN on test data using classifier trained with ColorSIFT feature model.

EXPERIMENTS AND RESULTS

In this section, we provide our experiments and results. First we present an overview of our Flickr dataset and gather statistics related to photos and their metadata, then our results of binary image classification. Based on these results, we are able to identify fire events across the globe, one event we use as a case study.

A significant proportion of our photos contain location information. Two forms of location data can be obtained through Flickr; geotags and geocodes. In our collection of 114,098 photos, 23,154 (20.3%) contain a geotag and 2,852 (2.5%) contain a geocode. Figure 4 shows the distribution of geotagged photos in our dataset across the globe, the majority of these posts originate from North America, Australia and the United Kingdom. As a comparison, a map of all geotagged images collected by Graham, Hale and Stephens (2011) is also displayed in Figure 4. Interestingly, our geographic distribution does not contain clusters of geotags in areas that are present in this map, such as areas throughout Asia and Europe. One explanation for this is that in building our dataset only English terms were used to gather images from Flickr and thus we would expect fewer posts from non-English speaking countries. Due to the subject matter in the images, our distribution of geotagged images may be also an overall indication of high-risk bushfire and wildfire areas throughout the globe.

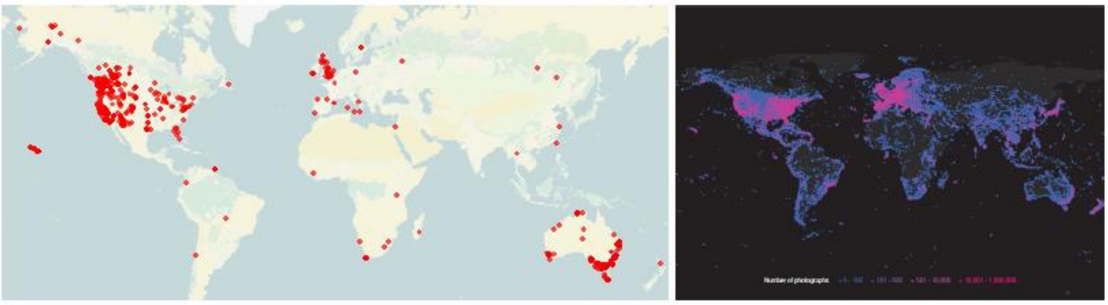

Figure 4. *Left:* Distribution of geotagged images in our dataset in 2015. *Right:* Map of all geotagged Flickr images on April, 2011. (*Source*[6]*:* Graham, M., Hale, S. A., & Stephens, M. (2011). *Geographies of the World's Knowledge.* London: Convoco.).

[6] Licensed under a Creative Commons Attribution-NonCommercial 3.0 Unported License http://creativecommons.org/licenses/by-nc/3.0/deed.en_US

Image Classification Results

After analysing our dataset, we then use our aforementioned image classifier to assign labels to images: either fire or not fire. As we are interested in images that can be mapped to a particular point in time and space, we run our fire classifier on only geotagged photos in our collection and find 1,925 matches for the fire concept.

Identification of Fires

With the images classified, we evaluate our method of filtering our dataset and gathering images related to fire events. The performance of our fire image classifier contained a high detection rate and relatively low false positive rate, therefore we use images classified as fire to predict the occurrence of fire events. Figure 5 shows a daily count of all geotagged images classified as fire in our dataset, and can be used to understand where major fires may have occurred by identifying peaks in activity.

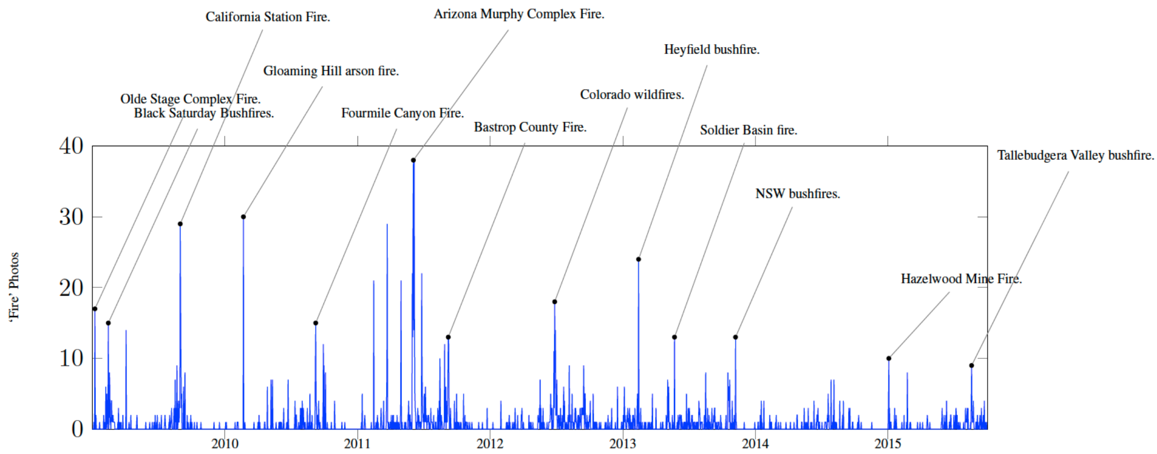

Figure 5. Daily count of all photos containing classified as fire, from January 1, 2009 to October 1, 2015. Large peaks are inspected to see what users have posted, and if they relate to a major fire the peak is labelled.

The number of images related to each increase in activity is not large (the most photos uploaded in a single day, classified as fire, was 38 images on June 4, 2011). Hence, we were able to manually assess the visual content of each image and its associated metadata. Below are major fires that occurred, as seen in Figure 5, which were also examined and confirmed with various sources:

- Olde Stage Complex Fire (Colorado, United States)[7] a fire that occurred on January 7, 2009.
- Black Saturday Bushfires (Victoria, Australia)[8] this fire began on February 7 and was contained on February 16.
- Station Fire (California, United States)[9] a fire spanning from August 26 to October 16.
- Fourmile Canyon Fire (Colorado, United States)[10] on September 6, 2010.
- Murphy Complex Fire (Idaho, United States)[11,12] was active between May 30, 2011 and was somewhat

[7] http://newsite.brfd.org/index.php/cms/the-olde-stage-fire-complex

[8] http://www.royalcommission.vic.gov.au/Finaldocuments/volume-1/PF/VBRC_Vol1_Chapter05_PF.pdf

[9] http://cdfdata.fire.ca.gov/incidents/incidents_details_info?incident_id=377

[10] http://earthobservatory.nasa.gov/IOTD/view.php?id=45675

[11] http://inciweb.nwcg.gov/incident/article/2268/11659/

contained on June 9, 2011.

- Bastrop County Fire (Texas, United States)[13] a fire on September 8 and September 9, 2011.
- Gloaming Hill arson fire (New Zealand)[14,15] confirm a fire caused by arson took place on February 21, 2010.
- Waldo Canyon Fire (Colorado, United States)[16] with a start date reported as June 23, 2012 and an estimated containment date of July 10, 2012.
- Heyfield bushfire (Victoria, Australia)[17,18] a fire occurred on 18 to 20 January 2013 however the images were uploaded on 13 February.
- Soldier Basin Fire (Arizona, United States)[19] started on May 17, 2013 and was contained on May 28, 2013.
- Blue Mountains bushfires (New South Wales, Australia)[20] that occurred beginning on October 16, 2013 and being declared contained on November 20, 2013.
- Hazelwood Mine Fire (Victoria, Australia)[21] most of the 8 images posted on January 3, 2015 referred to the fire in the previous year.
- Tallebudgera Valley bushfire (Queensland, Australia)[22] was occurring on August 19, 2015.

Some peaks did not correspond to major fires, for example a single-user may have uploaded and tagged many images as 'fire' and the image classifier incorrectly classified the images as fire.

Case Study of California Rim Fire

We use as a case study of the California Rim Fire, located in Stanislaus National Forest[23], as the images available allowed us to explore both spatial and temporal dimensions of the fire. It began on August 17, 2013 and the perimeter of the fire was contained on October 24, 2013. We restrict the photos and metadata we analyse to a specific time period (August 16, 2013 to November 30, 2013) and a geospatial region by using the bounding box shown in Figure 6.

---

[12] http://inciweb.nwcg.gov/incident/article/2268/11605/

[13] http://earthobservatory.nasa.gov/NaturalHazards/view.php?id=52045

[14] http://www.3news.co.nz/nznews/titahi-bay-blaze-suspected-arson-2010022117

[15] http://www.stuff.co.nz/national/crime/3356659/Teen-charged-over-Titahi-Bay-blaze

[16] http://inciweb.nwcg.gov/incident/2929/

[17] http://www.gippslandtimes.com.au/story/1244010/latest-map-of-the-heyfield-fire/

[18] http://www.abc.net.au/news/2013-01-18/bushfire-threatens-gippsland-towns/4470274

[19] http://inciweb.nwcg.gov/incident/3389/

[20] http://earthobservatory.nasa.gov/NaturalHazards/view.php?id=82211

[21] http://report.hazelwoodinquiry.vic.gov.au/wp-content/uploads/2014/08/Hazelwood_Mine_Inquiry_Report_Intro_PF.pdf

[22] https://newsroom.psba.qld.gov.au/Content/Local-News/14-South-East/Article/Tallebudgera-Valley-grass-fire-as-at-9am-Wed-19-Aug/1017/1058/8682

[23] http://inciweb.nwcg.gov/incident/3660/

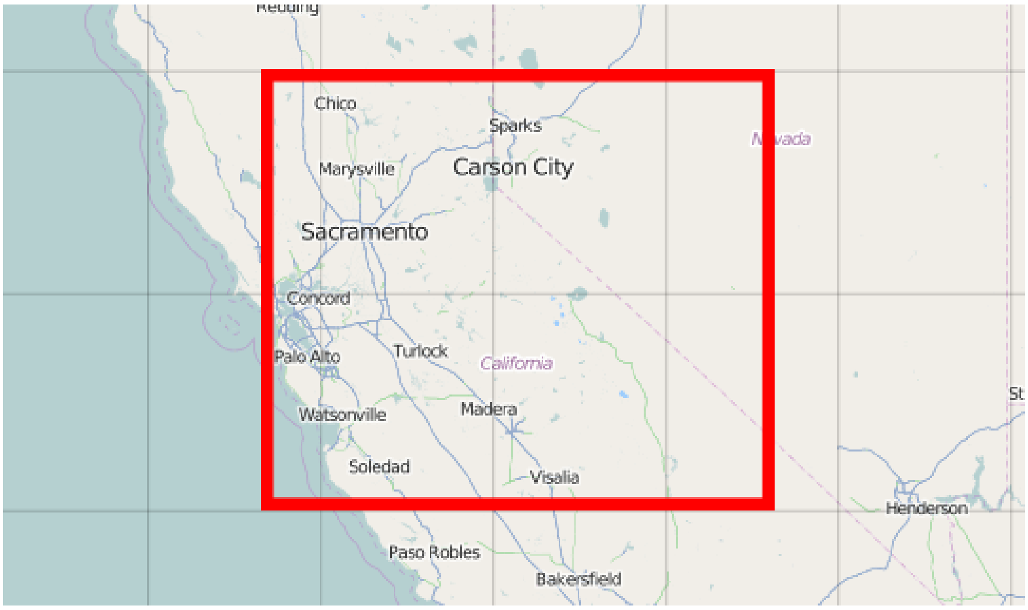

Figure 6. Bounding box used for our case study.

Figure 7. A sample of images classified: (a) image positively identified as containing fire [source: Rim Fire' by Steve Ryan (CC BY-SA 2.0)], (b) image falsely identified as containing fire [source: 'Reno Haze 2013' by Brian Ball (CC BY-NC-SA 2.0).

Visual Analysis

In total, 96 geotagged images were collected within our spatio-temporal restrictions. Of these 14 images were classified as 'fire' (two samples are shown in Figure 7). We observe that 9 of these images were correctly identified as containing fire or flames but 5 images are false positives – images that appear to contain no fire. From observation, an explanation as to why these images were classified wrongly is due to the color contained in the images. Almost all of them contain variations of red or orange and are similar in color to flames.

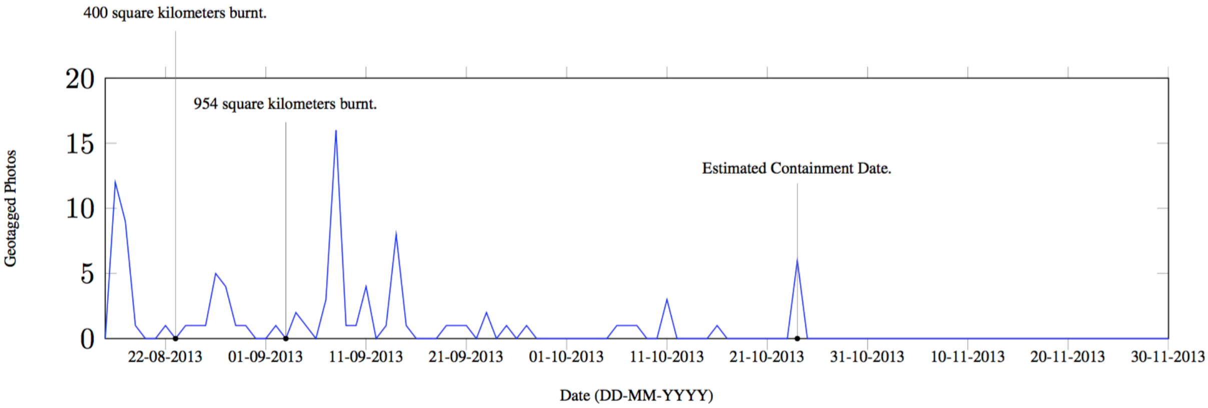

Figure 8. Chronological timeline of the California Rim Fire, with respect to the number of geotagged photos posted.

Temporal Analysis

For an analysis of the fire with respect to time, we create a timeline of the events and observe patterns in the data. Figure 8 shows a timeline of the fire, with respect to the number of geotagged Flickr photos posted each day. Ground truth data[24,25] is used to supply the amount of land burnt at particular points in time and is also labelled. There is an initial spike in the amount of photos posted the day after the fire first occurs, and again when the severity of the impact increases throughout the area. The amount of photos posted rises as the amount of land affected by the fire increases. On August 23, the fire had burnt through 400 square kilometres and steadily increased each day.

The most photos posted in a single day occurred on September 8, with 16 images posted. This occurs 2 days after it was reported that 954 square kilometres had been burnt. This data is keeping in line with the actual events transpiring in the area; as the fire escalated in size, there was an increase in the amount of photos shared on Flickr. As the fire was gradually becoming contained by local authorities and emergency services, the amount of posts evidently decreases. There are no geotagged posts within the area after October 25, shortly after the fire was contained. This sudden decrease in activity suggests that users are likely to remain active on photo-sharing sites during time periods when a fire is occurring, but lessen their presence upon the event's resolve. The last day that contains any photos is October 24, three days after the containment date.

Spatial Analysis

We use geotags to determine if a region contains fire-related activity or not, based on photos posted. A map containing all geotagged photos with our spatio-temporal restrictions is shown in Figure 9. For our spatial analysis, we use the location given by the USGS as our ground truth data[26]. The specific location reported in decimal degrees is 37.857, -120.086.

---

[24] http://earthobservatory.nasa.gov/IOTD/view.php?id=81971&eocn=image&eoci=related_image

[25] http://earthobservatory.nasa.gov/IOTD/view.php?id=81919

[26] http://landslides.usgs.gov/hazards/postfire_debrisflow/2013/20130817rim/

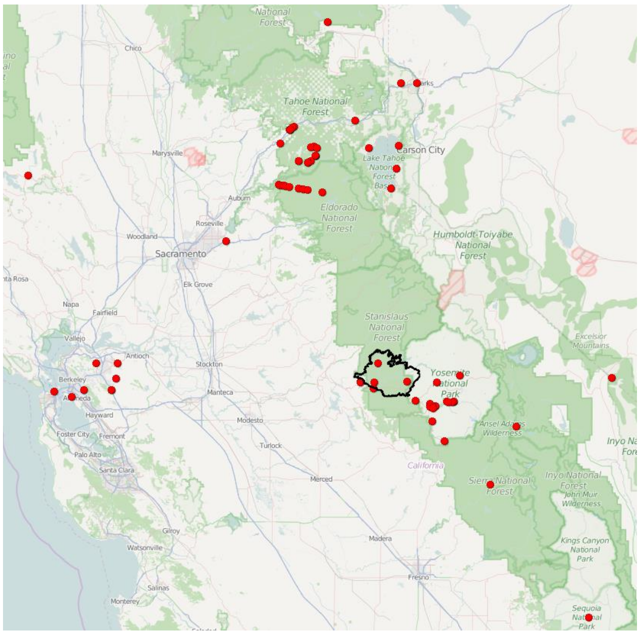

Figure 9. Map showing all geotagged photos within our bounding box. The approximate impact area of the fire is shown.

It is evident through visually inspecting our map of plotted geotags in Figure 9 that points group together in certain areas. We use the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm (Ester, Kriegel, Sander and Xu, 1996) to find cluster centers based on distance and density thresholds. Four clusters were identified, with distance of each cluster's central point from the fire listed in Table 2.

Cluster Number	Latitude, Longitude	Distance From Fire (Central Point)
#1	39.1283, -120.5248	146.221 km
#2	37.8112, -121.9931	167.956 km
#3	37.7460, -119.5333	50.21 km
#4	39.2750, -120.7090	166.508 km

Table 2. Cluster centers found after performing density-based clustering on our set of geotags.

An important attribute that geocodes provide is directional data. In EXIF tags, the tags *GPSImgDirection* and *GPSImgDirectionRef* [27] may be present, which indicate the direction the image was captured at. *GPSImgDirection* contains the values in degrees, ranging from 0.00 to 359.99. The tag *GPSImgDirectionRef* is supplied to indicate which bearing reference is used to give the direction of the image, with a value of 'T' or 'M' denoting True and Magnetic direction, respectively. Figure 10 shows two geocoded points found within the EXIF header data of individual images. The point shown on the left has a latitude of 39.270444 and longitude of -120.678833, with a direction of 126.0410765°T. The point shown on the right has a latitude of 37.817 and longitude of -120.109667, with a direction of 66.91381872°T. With the direction the image was taken at shown, it is evident that both images point near the approximate central location of the fire. We assume direction is reasonably accurate on this topographical map, as OpenStreetMap uses a spherical Mercator projection[28]. When more Flickr images have directional data in EXIF tags, then triangulation techniques may be used to help locate fires.

[27] http://www.exif.org/Exif2-2.PDF

[28] http://openstreetmapdata.com/info/projections

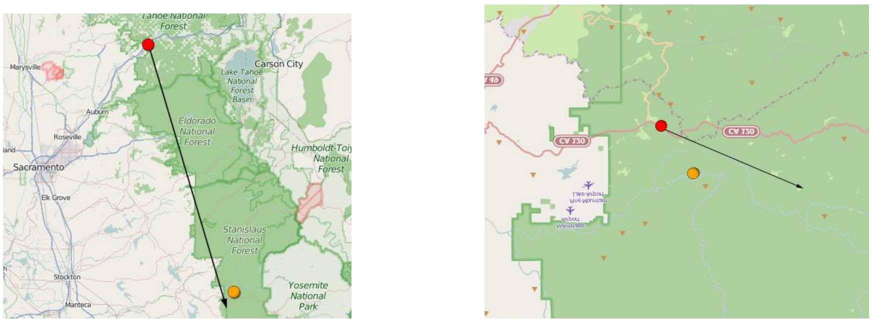

Figure 10. Geocoded points (red) obtained from image EXIF headers. The approximate direction the image was captured at is shown with arrows. The approximate central point of the fire (orange) is plotted as reference.

CONCLUSION AND FURTHER WORK

Our research proposed using a collection of photos gathered from social media to estimate the occurrence of a fire in a given time and place. We obtained a sum of 114,098 Flickr photos to accomplish this task. Of those photos collected, 20.3% were geotagged and 2.5% contained GPS data in the EXIF header.

We effectively applied image classification techniques to find images of 'fire' using ColorSIFT feature detection with the Bag of Features model. Using images assigned a 'fire' label, we analysed the temporal distribution of images in our collection with respect to the amount of 'fire' images posted. We were able to successfully identify fires that have occurred at a given time with confirmation from either dated ground truth satellite images or online authoritative sources. A case study of a major fire was then undertaken as an in-depth investigation, based on Flickr photos posted shortly before and after the fires occurred. An analysis of temporal metadata allowed us to see high periods of social media activity taking place during periods when there is an active fire. An analysis of geotags allowed us to estimate the location of the fire impact area using density-based clustering. In our case study of the California Rim Fire, we identified clusters located 50.21 km from the fire impact area.

We also performed a simple investigation of GPS data contained in EXIF headers. The results were very promising; directional data proved very useful in being able to gather a sense of where a fire occurred. While this metadata is much more difficult to extract and apply compared to geotags, the results show that the data contained in these headers may be useful. The ability to gather valuable directional data has a range of applications for geospatial tasks. Other elements contained in EXIF tags, such as the altitude the photo was captured at, should be explored.

We classified images in a collection of Flickr photos by considering only the visual content of the images. Elements such as text should be explored to gain additional information about photos. While we gather images related to a fire, some images may have been missed that contain important information located in the titles, tags and/or descriptions. Work should also include data from other social media networks as well to determine the variance in platforms. We use photos and metadata from Flickr, however this methodology can also be extended to other photo-sharing sites such as Twitter and Instagram.

REFERENCES

1. Abdel-Hakim, A. E. and Farag, A. A. (2006). CSIFT: A SIFT descriptor with color invariant characteristics. *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 2, pages 1978–1983. IEEE.
2. Amato, G., Falchi, F. and Bollitieri, P. (2010). Recognizing landmarks using automated classification techniques: evaluation of various visual features. *2010 Second International Conferences on Advances in Multimedia*.
3. Bay, H., Ess, A., Tuytelaars, T. and Van Gool, L. (2008). Speeded-up robust features (SURF). *Computer vision and image understanding*, 110(3): 346–359.
4. Cao, L., Luo, J., Kautz, H. and Huang, T. S. (2008). Annotating collections of photos using hierarchical event

and scene models. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1 – 8.

5. Çelik, T., Özkaramanli, H. and Demirel, H. (2007). Fire pixel classification using fuzzy logic and statistical color model. *Acoustics, Speech and Signal Processing, 2007. ICASSP 2007. IEEE International Conference on*, volume 1, pages I-1205-I-1208.
6. Chen, L. and Roy, A. (2009). Event detection from Flickr data through wavelet-based spatial analysis. *Proceedings of the 18th ACM conference on Information and knowledge management*, pages 523 – 532.
7. Chen, T., Lu, D., Kan, M.-Y. and Cui, P. (2013). Understanding and classifying image tweets. *Proceedings of the 21st ACM international conference on Multimedia*, pages 781 – 784.
8. Csurka, G., Dance, C., Fan, L., Willamowski, J. and Bray, C. (2004). Visual categorization with bags of keypoints. *Workshop on statistical learning in computer vision, ECCV*, volume 1, pages 1–2. Prague.
9. Ester, M., Kriegel, H.-P., Sander, J. and Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. *Proceedings of 2nd International Conference on Knowledge Discovery and Data Mining*, pages 226–231.
10. Fan, P., Men, A., Chen, M. and Yang, B. (2009). Color-SURF: A SURF descriptor with local kernel color histograms. *Network Infrastructure and Digital Content, 2009. IC-NIDC 2009. IEEE International Conference on*, pages 726–730. IEEE.
11. Firan, C. S., Georgescu, M., Nejdl, W. and Paiu, R. (2010). Bringing order to your photos: event-driven classification of Flickr images based on social knowledge. *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*, pages 189 – 198.
12. Graham, M., Hale, S. A. and Stephens, M. (2011). *Geographies of the world's knowledge*. London: Convoco.
13. Hadjisophocleous, G., Ouyang, J., Ding, G. and Liu, Z. (2010). *Study of video image fire detection systems for protection of large industrial applications and atria*. Carleton University, InnoSys Industries, Inc. and General Fire Technologies.
14. Huang, Y., Cheng, A., Hsie, L., Hsu, W. and Chang, K. (2011). Region-based landmark discovery by crowdsourcing geo-referenced photos. *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*, page 1141.
15. Li, L.-J. and Fei-Fei, L. (2007). What, where and who? Classifying events by scene and object recognition. *IEEE 11th International Conference on Computer Vision*, pages 1 – 8.
16. Li, Y., Crandall, D. and Huttenlocher, D. (2009). Landmark classification in large-scale image collections. *IEEE 12th International Conference on Computer Vision (ICCV)*, pages 1957 – 1964.
17. Lowe, D. (1999). Object recognition from local scale-invariant features. *Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 2.
18. McAuley, J. and Leskovec, J. (2012). Image labeling on a network: Using social-network metadata for image classification. *Proceedings of the 12th European conference on Computer Vision - Volume Part IV (ECCV'12),* pages 828 – 841.
19. Middleton, S. E., Middleton, L. and Modafferi, S. (2014). Real-time crisis mapping of natural disasters using social media. *IEEE Intelligent Systems*, 29: 9 – 17.
20. Muja, M. and Lowe, D. G. (2009). Fast approximate nearest neighbors with automatic algorithm configuration. In *International Conference on Computer Vision Theory and Application (VISSAPP'09),* pages 331–340. INSTICC Press.
21. Oliva, A. and Torralba, A. (2001). Modeling the shape of the scene: A holistic representation of the spatial envelope. *International Journal of Computer Vision* 42, pages 145 – 147.
22. Power, R., Robinson, B. and Ratcliffe, D. (2013). Finding fires with Twitter. *Proceedings of Australasian Language Technology Association Workshop*, pages 80 – 89.
23. Robinson, B., Power, R. and Cameron, M. (2013). A sensitive Twitter earthquake detector. *Proceedings of the 22nd International Conference on World Wide Web*, pages 999 – 1002.
24. Zhang, H., Korayem, M. and Crandall, D. J. (2012). Mining photo-sharing websites to study ecological phenomena. *Proceedings of the 21st international conference on World Wide Web*, pages 749 – 758.