PROCEEDINGS OF SPIE

SPIEDigitalLibrary.org/conference-proceedings-of-spie

Discovering anomalous events from urban informatics data

Kasthuri Jayarajah Vigneshwaran Subbaraju Dulanga Weerakoon Archan Misra La Thanh Tam Noel Athaide



Discovering Anomalous Events from Urban Informatics

Kasthuri Jayarajah[‡], Vigneshwaran Subbaraju[‡], Dulanga Weerakoon[†], Archan Misra[‡], La Thanh Tam[‡], and Noel Athaide[‡]

[‡]Singapore Management University, Singapore

[†]University of Moratuwa, Sri Lanka

ABSTRACT

Singapore's "smart city" agenda is driving the government to provide public access to a broader variety of urban informatics sources, such as images from traffic cameras and information about buses servicing different bus stops. Such informatics data serves as probes of evolving conditions at different spatiotemporal scales. This paper explores how such multi-modal informatics data can be used to establish the normal operating conditions at different city locations, and then apply appropriate outlier-based analysis techniques to identify anomalous events at these selected locations. We will introduce the overall architecture of sociophysical analytics, where such infrastructural data sources can be combined with social media analytics to not only detect such anomalous events, but also localize and explain them. Using the annual Formula-1 race as our candidate event, we demonstrate a key difference between the discriminative capabilities of different sensing modes: while social media streams provide discriminative signals during or prior to the occurrence of such an event, urban informatics data can often reveal patterns that have higher persistence, including before and after the event. In particular, we shall demonstrate how combining data from (i) publicly available Tweets, (ii) crowd levels aboard buses, and (iii) traffic cameras can help identify the Formula-1 driven anomalies, across different spatiotemporal boundaries.

Keywords: Multi-Modal Sensing, Information Fusion, Urban Analytics, Event Detection

1. INTRODUCTION

With the increased attention on developing technologies for smarter cities, we are beginning to see increased use of data from a diverse set of multi-modal sensors that are being deployed in various components of the urban infrastructure. Examples include cameras mounted along roads and highways (that help monitor realtime traffic congestion levels), and buses equipped with location sensors and contactless payment modes (that help monitor bus movement and occupancy levels). On the other hand, users of social media platforms such as Twitter and Instagram serve as a set of *distributed social sensors*, voluntarily sharing multimedia content related to events that occur in their localities.

Occupancy features extracted from such multi-modal, multi-source sensory information, can serve as accurate indicators of population levels at their respective locations. Longitudinal observations of such occupancy levels, or the level of crowdedness, can in turn help establish the normal operating conditions at those places, and establish how such conditions vary with temporal factors under normal conditions. For example, such variation is likely to depend on both the day of the week (e.g., the downtown areas could be expected to be less crowded during the weekends) and the hour of the day (e.g., the morning and evening rush hours are generally significantly more crowded than the remainder of the day). In previous work, researchers have successfully exploited such variations across a variety of observational modes (e.g., traffic flows on streets or call volumes on the cellular network) to isolate periods of anomalous behavior.

Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR VIII edited by Tien Pham, Michael A. Kolodny, Proc. of SPIE Vol. 10190, 101900F · © 2017 SPIE · CCC code: 0277-786X/17/\$18 · doi: 10.1117/12.2262404

[†] Work done during an internship at Singapore Management University. Further author information: (Send correspondence to A. M.)

A.M.: E-mail: archanm@smu.edu.sg, Telephone: +65 68085202

K. J.: kasthurij.2014@smu.edu.sg, V. S.: svignesh@gmail.com, D. W.: dulanga.weerakoon.lk@ieee.org, L. T.: tlathanh@smu.edu.sg, N. A.:noelathaide@smu.edu.sg

In this work, we focus on the problem of detecting anomalies or *events* at the city scale. In particular, we focus on a certain class of events that occur during a particular time and day but that potentially have a long, enduring impact or aftermath. Examples include: emergencies (such as gas leaks or building fires) that disrupt city traffic for hours afterwards, or preplanned events such as marathons or street festivals that can result in disturbances prior to, during and after the nominal period of the event.

We focus on the city-state Singapore whose government-spearheaded open data initiative has resulted in the public accessibility to transport-related APIs *. Among other attributes, such transportation-related data includes (i) images captured along major highways and expressways by traffic cameras; (ii) geocoordinates, expected arrival times and loading levels of public buses, and (iii) geocoordinates of available licensed taxis island-wide. As an alternative to such governmental data sources, we can also obtain trending topics (that may relate to specific physical events in the locality, identifiable by the keywords used) and geo-coordinates of Tweets using the publicly accessible Twitter API[†]. Given longitudinal observations, it is possible to then identify time periods and locations of anomalous activity using statistical outlier detection techniques, for each modality of sensing. More interestingly, besides simply detecting anomalies, social sensing data can be used to *explain* the cause of the anomaly as shown in previous works.¹

In this paper, we study the case of the Singapore Grand Prix, a major annual car racing event that happened in Singapore during the period of September 2016. The actual race happened on the 18th of the month, with practice races on the two days preceding that (16th and 17th) and a variety of entertainment events scheduled all three days. An important aspect of the race is that it takes place on streets that are otherwise part of the public road network in Singapore–to accommodate the special arrangements for the race, these streets are blocked completely from two days prior to the event to two days after (with the exception of certain times in the early hours of the morning). Consequently, although the race itself is a single day event, its effects persist for a significantly longer duration.

Our goal is to explore the extent to which we can detect the event, and localize it spatiotemporally (i.e., narrow down the duration and the locations where the race actually took place) using this combination of physical and social media sensors. We shall observe that although the Twitter stream is able to capture the actual race location and time accurately, it fails to capture the road block that persisted before and after the race day. On the other hand, the physically sensed data, comprising the street congestion levels (inferred from nearby traffic cameras) and the loading levels of buses in the vicinity of the race area, were able to capture the persistent anomalies. Interestingly, we shall see that the temporal scale at which detection is performed impacts its capability to detect the anomaly, depending on the source's spatial deployment resolution or the distance from the event's epicenter. Accordingly, different sensor feeds have different capabilities of spatiotemporal discrimination. Our work in this paper is part of a longer-term exploration of the question: *How can we improve the accuracy*, and *the spatiotemporal resolution*, *of such urban event detection*, *by harnessing these different discriminative capabilities of such multi-modal physical and social sensors?*

In this exploratory work, we make the following contributionas:

- We present and describe the architecture for socio-physical analytics, and describe the problem of detecting events from direct and passive sensing using both physical and social sensors and discuss its feasibility; in particular, we look at images from eight traffic cameras surrounding the downtown area of Singapore, 86 bus services servicing 116 bus stops servicing the area, and public Tweets posted by users from Singapore, all during the month of September, 2016.
- We employ a statistical outlier detection technique, the *Local Outlier Factor*, in detecting a large event and a temporary road block that was in effect before, during and after the event, using the congestion levels from the camera network, loading levels of the buses servicing the area, and the socially sensed information from the same period.

^{*}See Data.gov.sg, which provides access to searchable repositories of Singapore's public data. †https://dev.twitter.com/streaming/overview



Figure 1. A Framework for Socio-Physical Event Analytics.

• We provide empirical evidence into how different sensor modalities enable us to detect both the event and its longer-term impact. Such evidence provides us deeper understanding of how the insights from social and physical sensors may be fused in the future.

2. A FRAMEWORK FOR DETECTING CITY-SCALE ANOMALIES USING SOCIO-PHYSICAL ANALYTICS

In this section, we propose a socio-physical analytics framework for combining data from multiple social and physical data sources and discovering anomalous events. We describe further, the parts of the framework that this work addresses.

In Figure 1, we illustrate the framework. The social feeds can be textual or multimedia, whose relevancy and credibility are established, given their subjective nature, upon which *topics* are detected based on the presence of keywords. The details of this processing pipeline for social sensors is described in detail in Jayarajah et al.² In this work, we assume the keywords related to the events we study are known.

In processing the urban informatics data, we note the requirement for multiple forms of detecting occupancy. The reasons for this are two-fold: (1) first, the modality and structure of the data source are different (e.g., traffic camera returns images, while the bus stop data is available in a structured JSON format), and (2) the temporal resolution of the feeds from different sensors are often significantly different (e.g., most traffic cameras return an image roughly every 3 minutes while the bus occupancy and location information is updated once every minute). We describe how we bring data sources of differing resolution to a common form in Section 2.1 and how we extract occupancy levels from the different sources in Section 3.3.

The Normalcy Modelling component is optional based on the definition of the consequent stage (i.e., anomaly detection). In the case of supervised anomaly detection, this component extracts the patterns of the normal operating conditions, given that such a labelled training set exists. Approaches to do so can include statistical techniques (e.g., modeling underlying distributions) as well as machine learning based methods. When this is not the case, we rely on unsupervised approaches such as the *Local Outlier Factor* which we describe in Section 2.1. Thus, in this work, we do not assume any explicit human effort in labeling or annotating the physical sensor data.

Once anomalies are detected from independent sensor modalities, the Sensor Fusion component outputs a fused detection decision that takes into account the uncertainty in each independent source as well as the support across sources. In this work, we share a rudimentary approach for *localizing* the event (based only on the social data, see Section 2.2). In ongoing work, we are exploring extensions to this basic localization approach that leverage on the spatiotemporal propertiers of LOF-based outliers across multiple sensing modes. The Event Explanation component extracts root causes and semantics of the detected anomalies.

2.1 Anomaly Detection

We describe our methodology for detecting anomalies based on observations from a physical sensor node as a sequence of steps as follows:

- 1. Extract occupancy level: For some observation window t, we denote the observed occupancy levels as f_t . The average occupancy level is considered when multiple observations are available within t.
- 2. Extract temporal controls: For each observation, additional temporal controls such as the day of the week (dow) and the time of the day (bin) are extracted. In our evaluation, we consider 7 values for dow (one each for each day) and five variants of the bin (i.e., within 1, 2, 4, 6 and 8 hour(s) of the event time). In general, multi-hour bin values are likely to be more meaningful as they capture distinct periods of a typical day–e.g., evening rush hours.
- 3. Address Temporal Resolution: The choice of *bin* for each independent source is based on its respective sensitivity to an external disturbance (i.e., the anomalous event). Depending on the size of *bin*, the number of samples or observations of each source will differ across the sources. This will cause ambiguity at the fusion stage (not discussed in this paper). To overcome this, we propose the use of histograms (of fixed and matching lengths) of observation values as opposed to the raw time series of observations.
- 4. Outlier detection: Then, for each dow and bin, we compute the Local Outlier Factor (LOF) of each day of observations. The LOF is an unsupervised outlier detection technique that does not require explicit labelling of normal operating conditions. Typically, a LOF score > 1.2 is considered as indicative of an anomaly.

We consider each post as a tuple, < time, keywords, latitude, longitude >, where the latitude and longitude pair is available only for a subset of posts and the keywords contain at least one of the F1-related keywords. We first consider two count features, namely, the total number of event-related posts over an entire day and the total number of event-related tweets posted only between 7 PM and 9 PM, which is a 2-hour window centered around the event start time. Then, we compute the outlier score using the Local Outlier Factor algorithm, over all days.

2.2 Event Localization using Social Sensing

For the event-related tweets from each hour on the event day, we further subsample those tweets tagged with a geocoordinate. If N_h is the number of geotagged tweets from hour h, then we localize the event center as follows:

$$C_{h} = \begin{cases} t_{C,h} & N_{h} = 1\\ mean(t_{C,h}) & N_{h} = 2\\ centroid(maxCluster(t_{C,h})) & N_{h} > 2 \end{cases}$$
(1)

Here, $t_{C,h}$ is the set of geocoordinates from tweets for the hour h and maxCluster is the largest cluster returned by the k-means algorithm and centroid returns the centroid of that cluster. In our analysis, we take k = 2. If there are no geo-tagged tweets during an hour, then the event is not localized.

We define two measures, the quality of localization and its accuracy in order to evaluate the effectiveness of using geotags accompanying the event-related posts for fine-grained localization of the event.

Localization Quality: We measure the quality of the estimation based on the discriminating nature of the densest cluster. We define the quality, Q_h as:

$$Q_{h} = \begin{cases} \frac{|maxCluster(t_{C,h})|}{|nextCluster(t_{C,h})|} & N_{h} > 2\\ undefined & otherwise \end{cases}$$
(2)

The function nextCluster returns the second largest cluster and the operator |.| returns the magnitude or size of the returned cluster. The quality of the estimation basically determines how large or dense the largest

Proc. of SPIE Vol. 10190 101900F-4



Figure 2. Physical Locations of the Traffic Cameras and the Bus Stops Considered in this Study.

cluster is in comparison to the next largest cluster. If the two clusters have approximately the same size, it could mean that the event is being localized to two separate centers (if they centroids are sufficiently farther apart). We borrow this notion of quality from Giridhar et el.³

Localization Error: Next, we define the error in location estimation as the deviation from the known event center C_q . We define the error as:

$$E_h = haversine(C_q, C_h) \tag{3}$$

For each hour, we compute the Haversine distance between the known event center and the estimated event center C_h .

3. THE SINGAPORE GRAND PRIX – A CASE STUDY

In this section, we present a brief background of the Singapore Grand Prix (SGP) event (also referred to as the Singapore F1 Night Race), and the physical and social data we used in comparing the two sources with respect to their ability in detecting persistent effects of the event as well as the actual event.

3.1 Overview of the Event

The SGP race took place on the 18th of September in 2016, with the two preceding days hosting several practice and qualifying races. The 3-day period also showcased a variety of entertainment events. As the race circuit comprises roads that the public uses for commuting during normal days, the area was cordoned off from 14th September till 19th September to enable the installation (and subsequent dismantling) of race-related infrastructure. In Figure 2, the area marked by the black rectangle denotes the roads blocked during this period.

3.2 Socially-Sensed Data

We scraped posts on Twitter originating from Singapore during the period of September, 2016 using the publicly available API. Observing keyword trends is a well-studied methodology for detecting events in the social media data mining community. We use a set of keywords ({'formula1', 'sgp', 'racing', 'singaporef1', 'singaporegp', 'sgf1', 'sggp', 'formula1', 'nightrace', 'f1nightrace'}) commonly used by Twitterers during the F1 period to refer to the event to identify tweets related to the event. Further, we also extracted geocoordinates of where the post originated from, where possible, along with other metadata including the ID of the user and the time of the post.

In Table 1, we summarize the details of the dataset. In Figure 3, we plot the number of daily posts (left y-axis) against the number of posts related to the F1 event (right y-axis) for each day in September, 2016 (x-axis). It is clear that although the total number of posts decreased during the event period, the volume of posts related to the event spiked during the event period, specifically on the event date (i.e., 18th September 2016), validating our assumption that the keywords have utility in detecting the event.

Total number of posts	21,527,567
of posts with at least one keyword present	t 6,228
umber of posts with geo coordinates	208,349
Unique users	2,693,168
Table 1. Basic information of the dataset.	
T (() ()	
Traffic Camera	Bus Services
September 2016	Aug - Dec 2016
70 cameras deployed along	5000 bus stops servicing 331 routes
expressways, gateways and	
checkpoints	
1 - 15 mins	$1 \min$
	Timestamp, Estimated Arrival
	Time, loading information and
	GPS coordinates of the
Timestamp, Camera ID, Tramc	immediate and subsequent bus
Image	for each serviced route of the
	immediate and subsequent bus
	for each serviced route
	Total number of posts of posts with at least one keyword presen umber of posts with geo coordinates Unique users Table 1. Basic information of the datase Traffic Camera September 2016 70 cameras deployed along expressways, gateways and checkpoints 1 - 15 mins Timestamp, Camera ID, Traffic Image

Table 2. Dataset Description. The data is publicly available through the DataMall, an open data initiative from the Land Transport Authority of Singapore

3.3 Physically-Sensed Data

In this work, we consider two physical sensing sources, traffic camera feeds and bus stop data. Both sources differ in spatial resolution, update frequencies and data structures (see Table 2).

For our analyses, we consider the image feed from the eight cameras closest to the downtown area, where the Singapore Grand Prix was held. As the image update rate varies between 1 - 15 minutes across the cameras and hence, we fixed the observation interval to be 15 minutes. Over the month of September, 2016, we collected a total of 109,962 images and processed each image to extract the occupancy/congestion level as outlined below.



Figure 3. Daily number of total posts and posts with keywords related to the Singapore Grand Prix.

- 1. Defining Region of Interest (ROI): For each of the cameras, we picked an exemplary image and defined the portions corresponding to the navigable part of 'roads" and performed clustering on the road pixels. Processing only the ROI instead of processing the entire image improves processing time and reduces the possibility of false positive object detections.
- 2. Corner detection: Next, we detected the Harris corners within the road segments. These corners represent the corners of vehicles.
- 3. Image segmentation: Once the corners are detected, we used the watershed algorithm (used for image segmentation) to cluster the corners into segmented vehicles. As an artefact of this process, it is possible that multiple vehicles which are back to back (or partially occluded) can be identified as one single vehicle.
- 4. Congestion level estimation: Finally, we count the number of pixels, normalized by the pixel count of the ROI, that are part of the segmented vehicles within the ROI. This pixel counting approach is preferred over an alternative approach of explicitly counting the number of distinct vehicles, as it is more robust to the overlap and occlusion artefacts in the images.

For corner detection and watershed image segmentation, we used the implementations available in OpenCV. To understand the accuracy of this vehicle detection process, we conducted two benchmarking evaluations: (1) first, we used the MIT Car dataset [‡], and (2) 60-70 images from each of the 8 traffic cameras, randomly sampled. For each dataset, we marked ground truth as the number of vehicles seen in the image and the detection results as the number of vehicles contained within the "vehicle segments". We achieved an accuracy of 79.75% and 90% for the MIT and our datasets, respectively.

Next, for the bus stop information available, we extracted the average loading level across the next arriving bus of each bus service that serves that bus stop, during the observation interval. The loading level is one of four values: not available (0), limited availability (1), standing available (2) and seating available (3).

We defer the consideration of more complex features, such as the inter-bus stop speed distribution or the wait time at an individual bus stop, to future work.

4. PRELIMINARY RESULTS

In this section, we investigate two specific questions: (1) are there any fundamental differences in the feasibility of detecting anomalous events and their persistent effects (before as well as after the event) across the sources, and (2) the ability of the sources (in this work, the Tweets) to localize the event.

4.1 Anomaly Detection

Here, we report on the outlier scores computed for each data source, separately.

4.1.1 Detection using Traffic Congestion Levels

In Figure 4, we plot the outlier scores of each camera on the SGP race day for different *binss*. As expected, the highest score is seen for the camera closest to the event area right around the race time (*Camera*3798, 8 PM to 10 PM) indicative of traffic coming into the downtown area to watch the race. On the other hand, camera 1502 which is at the bottom highway (in Figure 2), has detected anomalous traffic over the longest *bin*, likely representative of traffic that shifted away from the downtown area to bypass the road blocks.

Further, in Figure 5, we plot the computed LOF score for each day in September, 2016, for camera 3798. The x- axis depicts the count of days since the actual race date. We observe that the camera was able to detect the race date and the two practice dates prior to that, but not the days where there was only a road block on normal streets in effect. This is expected as the camera oversees the congestion level on the highway which may not have been affected considerably by the road blocks.

[‡]http://cbcl.mit.edu/projects/cbcl/software-datasets/CarData1Readme.html



Figure 4. The LOF scores of each camera on the race day for different definitions of *bin*. Camera3798, which is located near the exit towards downtown along the *East Coast Parkway*, shows anomalous traffic on the F1 race day.



Figure 5. The computed LOF score for each day in September, 2016, for Camera 3798. The actual event date is marked as Day 0.

4.1.2 Detection using Crowd Levels in Buses

Similar to the case of traffic cameras, bus stops on roads that are part of the Marina Bay Circuit, show anomalous behavior across the entire period (from 14th September to 19th September) due to the road block, thus capturing both the actual event plus the persistent effects. In addition, bus stops along roads in the vicinity of the event were also effected. For example, in Figure 6, we plot the outlier scores of bus stops along the *Victoria Street* which shows anomalous behavior for certain stops, particularly around the event start-end period.

4.1.3 Detection using Tweets

In Figure 7, we plot outlier scores based on the all day count (blue, solid line) and the 2 hour count (red dashed line) for each day centred around the event day (marked as 0) for the month of September 2016. The black horizontal line represents an outlier score of 1.2 - an outlier score below this value is typically considered as normal.

Contrary to the case of the physical sensors, we observe that the whole-day count of Tweets containing SGP-related keywords exhibits an anomaly only a day prior to event and on the actual event. If we consider only the count of Tweets during the start hour of the event (bin = 8 - 9 PM), then the detection capability improves - the event is detected two days prior and till a day later. During the days leading to the actual race, there were a number of other events which we believe may have caused an increase in the number of tweets related to the event even before it actually took place – for instance, there were two practice races and a qualifying race held during the evenings. Additionally, we also observe that there were an anomalous number of posts related to the event on two other days – upon closer investigation, we observed that these dates coincided with the dates the organizers ran marketing campaigns or posts for promoting the event.

Next, we investigate the ability to detect the event as the event day unraveled. To this end, we computed the LOF score for each hour of the event day using the count of the event-related tweets for that hour. For each hour of the day, there were 30 data points (for each day in the month), for which LOF was used to detect outliers present. In Figure 9, we plot the outlier score of each hour (blue, solid line) for each hour of the day centered around the race start time – i.e., the start time 8 PM is represented by the 0th hour. Interestingly, we observe that the highest outlier scores are observed for the two hour preceding the actual start time – this is in agreement with the anticipatory behavior of spectators where they gather at the event venue ahead of time, both to avoid the rush close to the start time as well as to share their excitement with the crowd.

Our results reveal an important observation: the peaks of social sensing seem to act as a biased estimator, occuring before or just around the start of an event, and have sharper peaks. In contrast, the outliers for physical sensing (both via traffic cameras and bus occupancy levels) are more unbiased (centralized around the event's duration), but persist over a longer time period. Accordingly, our proposed multi-modal fusion algorithms will have to be developed to explicitly handle the different bias and spread of each such estimator.

4.2 Anomaly Localization

As described in Section 2.2, we evaluate the quality and accuracy of the localization possible with a single source – i.e., tweets in this work. In Figure 8, we plot the quality for each hour on the race day, and in Figure 9, we plot the error in localization in meters, for each hour. For the error calculation, we consider the known epicenter of the event as the center of the Marina Bay Circuit[§]. We make the following observations.

1. The quality of localization becomes highest during the two hours prior to the actual race start time, complementing our observations of the outlier score.

[§]https://en.wikipedia.org/wiki/Marina_Bay_Street_Circuit



Figure 6. The LOF scores of the bus stops along Victoria Street which is one of the roads adjoining the race circuit on the F1 race day.



Days

Figure 7. The LOF scores for each day in September, 2016 with the event day represented as the 0th day. The blue solid line represents the outlier score based on the total count of event-related tweets over the entire day. The red dashed line represents the total count of event related tweets in the two hour period centered around the event start time- i.e., 7 PM to 9 PM.

2. During the period before the actual start time and the event duration, we observe that the localization error is in the 500-600 meter range. Considering the fact that the Marina Bay Circuit itself is 5 km long, a 500 meter error can be considered reasonable.

4.3 Summary

In summary, we observe that the physical sources, the bus loading level, in particular in this case, is better at detecting the sustained impact of the effect on normal mobility, as compared to the socially sensed data. For each of the sources, we compute the *recall* of event detection as the fraction of days the source detected as anomalous out of the 6-day extended event period between 14th September and 19th September.



Hours

Figure 8. The outlier score for each hour on the race day and the quality of localization during that hour.



Hours

Figure 9. The outlier score for each hour on the race day and the distance error in localizing the event during that hour.



Figure 10. The recall of detection for each of the sensor sources considered and the actual or deduced distance of the sensor from the assumed event center.

In Figure 10, we plot the recall of bus stops along streets that were part of the Marina Bay Street Circuit (red circles), bus stops that were along streets immediately next to the circuit (blue triangles), recall of camera 3798 (orange plus) and the recall of tweets (purple cross). The x - axis shows the distance of the sensor (in the case of physical sensors) and the estimated location (in the case of Tweets) from the assumed ground-truth location of the circuit. It is clear that a majority of the bus stops were able to detect all days affected (*recall* = 1) whilst the camera and the social sensors were observed to be less sensitive to the road blocks.

5. DISCUSSION

In this section, we summarize some of our key limitations in the current study and our plans for future work.

A key limitation in the current study is that we consider the observations from each sensor node (of the same modality) as an independent source. While this is useful in capturing large deviations, this may not capture small, yet meaningful deviations, as seen in the case of traffic cameras. We believe that utilizing the underlying network structure (arising out of the road topology and bus routes) of the sensor nodes will allow us to better understand small changes to the occupancy/traffic levels, as well as shifts in traffic from one route to the other. Another key limitation we identify is the loss of accuracy in detecting images (and hence, extracting congestion levels) under poor lighting conditions. We observe that the accuracy of detecting vehicles drops

drastically during the evenings (yellow traffic lamps) and late night/wee hours where the infrared mode of the cameras are triggered. With the rapid advances in deep-learning based object detection and larger amounts of training, we believe that the accuracy could be improved in future.

Our ongoing and future work is expanding our investigations along several distinct dimensions:

- 1. Exploit Topological Network Structure: As explained, we shall not only look at the LOF score of each sensor node in isolation, but look at the *correlation* among neighboring sensors as an important observation metric, and derive LOF scores over such correlation scores. Such correlation-aware outlier detection patterns can not only improve the spatiotemporal localization of an individual event, but is also likely to help in situations where multiple distinct events, with possibly overlapping zones of impact, contribute to the anomalous observations.
- 2. Consideration of Additional Sensor Modalities: We are also working to expand both the number of distinct sensor modes (e.g., taxi availability, car park occupancy) and the set of features we employ in proxying occupancy (e.g., estimated bus arrival times). Our goal is to first build up a mode-specific understanding of the bias and spread properties of these estimators, and to also determine if these properties are affected (for a single sensing modality) by the type of event–e.g., does a crowd's Tweeting pattern differ accross sports, political and emergency response events?
- 3. Fusion of Multi-Modal Sensor Information: Eventually, we hope to move beyond mode-specific anomaly detection and localization to a more comprehensive framework, where we fuse the distinct information from the different sensing modes. The vision is to develop a detection-cum-estimation framework that accommodates, and in fact capitalizes on, the different spatiotemporal bias and spread properties of each mode-specific estimators. One additional intriguing alternative to such statistical estimation techniques (which require manual crafting of the features and the fusion logic) is to explore a deep-learning based approach that operates on such multi-dimensional, multi-modal data.

6. RELATED WORK

We briefly describe related work in the subdomains of urban event detection, event detection in the social media, anomaly detection in transport (particularly in road networks), and review some early works in socio-physical analytics.

Urban Event Detection: Recently, Konishi et al.⁴ discuss an approach to *predicting* irregularities (e.g., large scale events) ahead of time using a two-step modeling process. They assume the availability of a transit App used by mobile users where they are able to query route information by specifying a origin location, destination location and a time in future. First, using route search logs (for future and current times), the authors model a short term population model using auto-regression and a long term, routine population model using a bilinear Poisson regression with a 14-dimensional external factor (7 days X {holiday, working day}). Through the use of a deviation measure and a corresponding threshold, they are able to *predict* the occurrence of anomalies, or events, ahead of time. Further, in Jayarajah et al.⁵ and Nayak et al.,⁶ detection of anomalous events using occupancy levels, as well as changes to network properties, has been explored and demonstrated both at a campus level (on indoor movement of students) and at city-scale (on contactless payment card taps at train stations).

Event Detection using Social Media: Since the work of Sakaki et al.⁷ on detecting and tracking Earthquakes from user posted information on Twitter, the possibility of detecting and monitoring events using Social Media has been studied extensively. Typically, social media-based event detection is useful for identifying events with at least moderately high impact (e.g., train network outages), as the overall fraction of such geographically-tagged, event-specific posts is usually low. We have recently focused on utilizing multi-modal social data (text, images and metadata) for detecting and characterizing urban events. In particular, in Jayarajah et al.² we describe a framework for the fusion of multimodal social data for the detection and classification of events and the determination of anomalies. Also, in Jayarajah et al.,⁸ we describe how, the availability of geo-tagged Instagram posts can be utilized in disecting large, urban events to further characterize them as a series of meaningful, micro-events.

Anomaly Detection in Transportation Networks: Previous works on anomaly detection in transportation have looked at varied aspects of the problem including the detection of anomalies, understanding the saptiotemporal ordering and finding the root causes of the anomalies. Pang et al.⁹ detect contiguous, spatiotemporal cells as anamolous regions using Likelihood Ratio Tests and porpose a pruning strategy to improve the scalability of the approach. Further, Liu et al.¹⁰ proposed a formulation for "causal outlier detection" for detecting the emergence, propagation and disappearance of outliers (e.g., traffic jam) in addition to detecting them. They represent the road network and flows between them as a graph of regions and detect spatio-temporal using the minimum distortion measure. More importantly, they infer the causal relationship of the detected anomalies using a frequent subtree algorithm that is loosely based on association rule mining. Further still, Chawla et al.¹¹ attempt to identify routes in a road network that has caused anomalous traffic in the regions of interest taking a 2-step approach: (1) first they detect anomalous links using Principal Component Analysis (as seen in many works on network traffic anomaly detection) and (2) using a link-route matrix, they detect which routes were root causes for the detected anomalies using L1 machinery.

Socio-Physical Analytics: Since recently, researchers have focused on understanding the challenges and opportunities in combining physically sensed data with socially sensed data. Gao et al.¹² were one of the first to recognize the availability of multiple, independent data streams on the web and the need for combining them for situation recognition and awareness. They describe a data format, *E-mage*, which helps establish commonality in terms of space and time (STT) with each disparate stream, at its own spatiotemporal granularity, will measure a "theme" (a numerical value). They describe a framework which they implement and demonstrate applications. The framework includes a data ingestor (for taking in continuous data streams and a data wrapper to convert it to STT observations) and a set of operators for operation over time windows, for situation awareness. Recently, Misra et al.¹³ describe a vision for socio-physical analytics and identify the key challenges that need to be addressed in realising this vision; this includes (1) fusing multi-modal, multi-timescale information, (2) describing a multi-layer network of varied modes of sensors and understanding its properties, and (3) accounting for complementary and conflicting information across the modes. Wang et al.¹⁴ were one of the first to look at the fusion of GPS probes from vehicles and socially sensed data. The work tries to solve the data sparsity problem that is prevalent if one would only rely on a unimodal, sparse sensor (e.g., GPS). They attempt to *complete* a sparse congestion matrix using additional input from Tweets from the same period of time. More recently, Giridhar et al.¹ make the case for identifying root causes for sensor anomalies using social media data. They describe a system called "ClariSense" which detects anomalies in the physical sensor data and social data independently, and by matching them by space and time, tries to *explain* the physical sensor anomalies using the semantics learned from Twitter data.

7. CONCLUSION

In this work, we have explored the concept of using socio-physical analytics, i.e., the analysis of data provided by both physical and social sensors, to *detect* and *localize* urban events. We specifically focused on a moderatescale event, the annual Singapore F1 street race that results in changes in transportation flows over a period of 5-6 days. We showed that (i) traffic camera-based congestion estimates and (ii) public bus occupancy levels both reveal anomalies caused by this event, and that a LoF-based algorithm can identify the duration of such event-driven disruptions at both short and medium timescales. Additionally, we showed that appropriately filtered Twitter feeds could also help detect this event, even though the social media-based anomalies occurred primarily prior to the start of the event. Moreover, localization based purely on such Twitter feeds was able to achieve a spatial resolution of approx. 500 meters. Our work provides evidence that fusing such physical and sensor data can help both detect and localize such urban anomalies with high accuracy and fine granularity, provided we appropriately factor in the different bias and spread characteristics of each sensing mode.

ACKNOWLEDGMENTS

This material is based on research sponsored in part by by the Singapore National Research Foundation under its International Research Centre @ Singapore Funding Initiative, and in part by the US Air Force Research Laboratory, under agreement number FA2386-141-0002. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory or the U.S. Government.

REFERENCES

- Giridhar, P., Amin, M. T., Abdelzaher, T., Wang, D., Kaplan, L., George, J., and Ganti, R., "Clarisense+: An enhanced traffic anomaly explanation service using social network feeds," *Pervasive and Mobile Computing* 33, 140 – 155 (2016).
- [2] Jayarajah, K., Yao, S., Mutharaju, R., Misra, A., De Mel, G., Skipper, J., Abdelzaher, T., and Kolodny, M., "Social signal processing for real-time situational understanding: A vision and approach," in [Mobile Ad Hoc and Sensor Systems (MASS), 2015 IEEE 12th International Conference on], 627–632, IEEE (2015).
- [3] Giridhar, P., Wang, S., Abdelzaher, T., Kaplan, L., George, J., and Ganti, R., "On localizing urban events with instagram," in [Computer Communications (INFOCOM), 2017 International Conference on], IEEE.
- [4] Konishi, T., Maruyama, M., Tsubouchi, K., and Shimosaka, M., "Cityprophet: City-scale irregularity prediction using transit app logs," in [Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing], UbiComp '16.
- [5] Jayarajah, K., Misra, A., Ruan, X.-W., and Lim, E.-P., "Event detection: Exploiting socio-physical interactions in physical spaces," in [Advances in Social Networks Analysis and Mining (ASONAM), 2015 IEEE/ACM International Conference on], 508–513, IEEE (2015).
- [6] Nayak, S., Misra, A., Jayarajah, K., Prasetyo, P. K., and Lim, E.-p., "Exploring discriminative features for anomaly detection in public spaces," in [SPIE Defense+ Security], 946403–946403, International Society for Optics and Photonics (2015).
- [7] Sakaki, T., Okazaki, M., and Matsuo, Y., "Earthquake shakes twitter users: Real-time event detection by social sensors," in [Proceedings of the 19th International Conference on World Wide Web], WWW '10.
- [8] Jayarajah, K. and Misra, A., "Can instagram posts help characterize urban micro-events?," in [Information Fusion (FUSION), 2016 19th International Conference on], 130–137, IEEE (2016).
- [9] Pang, L. X., Chawla, S., Liu, W., and Zheng, Y., "On mining anomalous patterns in road traffic streams," in [Proceedings of the 7th International Conference on Advanced Data Mining and Applications - Volume Part II], ADMA'11.
- [10] Liu, W., Zheng, Y., Chawla, S., Yuan, J., and Xing, X., "Discovering spatio-temporal causal interactions in traffic data streams," in [Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining], KDD '11.
- [11] Chawla, S., Zheng, Y., and Hu, J., "Inferring the root cause in road traffic anomalies," in [Proceedings of the 2012 IEEE 12th International Conference on Data Mining], ICDM '12.
- [12] Gao, M., Singh, V. K., and Jain, R., "Eventshop: From heterogeneous web streams to personalized situation detection and control," in [Proceedings of the 4th Annual ACM Web Science Conference], WebSci '12.
- [13] Misra, A., Jayarajah, K., Nayak, S., Prasetyo, P. K., and Lim, E.-p., "Socio-physical analytics: challenges & opportunities," in [Proceedings of the 2014 workshop on physical analytics], 19–24, ACM (2014).
- [14] Wang, S., He, L., Stenneth, L., Yu, P. S., Li, Z., and Huang, Z., "Estimating urban traffic congestions with multi-sourced data," in [2016 17th IEEE International Conference on Mobile Data Management (MDM)], (2016).