Spatio-temporal Event Detection using Poisson Model and Quad-tree on Geotagged Social Media

Yasmeen George*†, Shanika Karunasekera†, Aaron Harwood†, and Kwan Hui Lim‡
*IBM Research - Australia
†School of Computing and Information Systems, The University of Melbourne, Australia
‡Information Systems Technology and Design Pillar, Singapore University of Technology and Design, Singapore
Email: georgey@ibm.com, {karus,aharwood}@unimelb.edu.au, kwanhui_lim@sutd.edu.sg

***Abstract*—Identifying events happening in a specific locality is important as an early warning for accidents, protests, elections or breaking news. However, this location-specific event detection is challenging as the locations and types of events are not known beforehand. To address this problem, we propose an online spatio-temporal event detection system using social media that is able to detect events at different time and space resolutions. First, we exploit a quad-tree method to split the geographical space into multiscale regions based on the density of social media data. Then, we implement a statistical unsupervised approach using Poisson distribution and a smoothing method for highlighting regions with unexpected density of social posts. Further, event duration is estimated by merging events happening in the same region at consecutive time intervals. A post processing stage is introduced to filter out events that are spam, fake or wrong. Finally, we incorporate simple semantics by using social media entities to assess the integrity, and accuracy of detected events. The proposed method is evaluated using Twitter and Flickr for the city of Melbourne based on recall and precision measures. We also propose a new quality measure named strength index, which automatically measures how accurate the reported event is.**

***Index Terms*—Online Event Detection; Social Media; Quad-tree; Poisson Distribution, Twitter, Flickr**

I. Introduction

Social media services, such as Twitter, are frequently used as a source of news and other information. Social media serves as an efficient source of breaking news compared to traditional media, which are either slow to pick up such information or do not give a complete and accurate picture of the news and events. Due to these reasons, many researchers and organizations are relying on social media for obtaining timely news. One emerging use case of significant importance is where social media information is used for real-time event detection. For example, governments and organizations may be interested in events that are occurring in a particular geographical area, such as detecting a bush fire near residential areas, traffic congestion and accidents on highways, protests and other security incidents in the city. Being able to promptly detect such events is important as this early detection allows the relevant authorities and organizations to make the necessary responses to address these potentially adversarial events.

The traditional approach to detect events in social streams is to track the aggregate trend changes based on the count of geotagged social media data at given location and time. This approach is very closely related to topic detection and tracking, where an event is conventionally represented by a number of keywords, topics or tweets showing bursts in appearance count, i.e. keywords that are mentioned significantly more often during a (not too short) time period than in the period preceding it [1], [2]. Most of these existing approaches detect events at fixed spatial and temporal resolutions, e.g., grids, which do not adequately capture the dynamic changes in tweeting volume across different areas and time [3]. However, real-life events can occur at any spatial or temporal resolution, which is not known a priori and, therefore, algorithms that have fixed resolution result in suboptimal performance. While there are some approaches that are designed for detecting events at multiscale spatio-temporal resolutions [4]–[6], they are essentially batch-based algorithms which are not directly applicable in the online real-time event detection scenarios. While there are a few online spatial event detection algorithms proposed in the literature [7]–[9], they are fixed in terms of spatial and temporal resolution. Many of these works also utilized a supervised approach to event detection which may not work well for new types of events.

Research Objectives and Contributions. Our aim is to detect spatio-temporal events from social media. It could be any event which is being discussed loudly (frequently) in a specific local or global area. The first challenge associated with event detection is that there is no consensus among researchers on the definition of an event. The second challenge is that the location, time and the scale of the events (both in time and space) are not known before hand. Finally, we are interested in detecting events in real-time from high velocity data streams and therefore, algorithms need to be single pass and computationally efficient.

We propose a novel approach to online spatio-temporal event detection that utilizes: (i) a quad-tree and Poisson model variant to dynamically identify events across different spatial scales; and (ii) a smoothing and filtering approach to

978-1-7281-0858-2/19/$31.00

effectively detect events with different temporal resolutions. The contribution of this paper can be summarized as follows:

- leveraging the quad-tree data structure for multiscale event detection;
- combining a Poisson model with a smoothing function for unsupervised event detection;
- a new event validation measure, strength index (SI), which automatically assesses the accuracy of the detected events by using social media entities; and
- evaluation of the proposed method using different social media datasets: Twitter and Flickr.

Structure and Organization. This paper is structured as follows. Section II reviews related work in the area of event detection, while Section III introduces the formulation of the event detection problem. Section IV describes our proposed algorithm for location-based event detection, and Section VI shows the experimental results of our proposed algorithm against various baselines. Finally, Section VII summarizes and concludes this paper.

II. Related Work

Many approaches have been developed for event detection in the spatial, textual (i.e. semantic) and temporal context, separately [3], [10]. But only few attempts are reported which combine spatio-temporal information for event detection. In [11], spatio-temporal events are detected by clustering the geotagged tweets, followed by topic modelling using the summarized words in each estimated cluster. Similarly, [12] adopted an approach of identifying topics associated with specific locations by applying Latent Dirichlet allocation on tweets posted in the same locality. Others like [13] combine clustering techniques with embeddings of tweet location, time, and text for event detection. The fact that not all of the tweets are geotagged restricts the accuracy of spatial-based event detection approaches. Although Twitter enables users to post tweets with their current locations (longitude and latitude), only an average rate of 0.85–3% tweets being geotagged per day, around 7,000,000 geotagged tweets are posted per day [14].

Another major challenge is the method of partitioning the geographical area. In this regard, using a uniform-grid, which applies an equi-width grid of a certain size over the data domain does not solve the problem for various reasons. First, a good method for choosing the grid size is required, which has not been accurately covered in the literature [15]. Second, fixed grid cells might not help in finding both local and global events. For example, using a low resolution grid for spatial data might capture only the global events occurring on the state or the country level, while a high resolution grid will detect the events on smaller scales (local events), i.e. within the community or the city where the grid cell ranges from 1km-50km. Another solution is to manually select a set of points of interests (POIs), where each POI is a fixed size grid cell. Following this approach, we can control the number of POIs based on tweets distribution density. For instance, areas in city centre might have many POIs with small grid cells, while areas far from the city might have few POIs with large grid size. But, having fixed POIs limits the location of detected events to the chosen POIs only. In addition, the manually selected POIs has to be done for each geographical area of analysis.

TABLE I
High level notation

p_x	A social media post p_x
t_x	Timestamp of post p_x
l_x	Location of post p_x
f_x	Features of post p_x
$\mathcal{S}$	A data stream of posts as a set, $\mathcal{S} = \{p_1, p_2, \ldots, p_n\}$
$\mathcal{W}$	A window (e.g. of size m), $\mathcal{W} = \{p_{n-m+1}, \ldots, p_n\} \subset \mathcal{S}$
$\mathcal{E}$	A set of posts $\mathcal{E} \subseteq \mathcal{W}$ that represents a spatio-temporal event

There are also many related work that study the general problem of event detection using social media, without an explicit focus on the spatial aspect of events. For example, numerous researchers have examined the problem of identifying trending and bursty events [16]–[18], and detecting controversial events [19], [20]. Weng and Lee [21] utilize (tweet) word signals derived from wavelet analysis, which are clustered together using a modularity-based graph partitioning to represent detected events. However, many of these works aim to detect events without considering the spatial aspects of these events. Others like Sakaki et al. [22], [23] first use a trained Support Vector Machine to determine if tweets are earthquake-related or not, then applies Kalman filtering and particle filtering on tweets to estimate the centres of detected earthquakes. Similarly, there are also various web and mobile applications for tracking general events or retrieving tweets related to specific events [24]–[26]. For a more detailed survey on general event detection, we refer readers to [27].

III. Problem Statement

In this section, we first introduce some basic notation and definitions used in our work, before formally defining the problem of spatio-temporal event detection.

A. Preliminaries

Table I summarizes the key notations used in our work, which we elaborate next.

Definition 1 (Social Media Post): We represent each social media post as $p = \langle t, l, f \rangle$, where each social media post p is associated with a timestamp t, location l and features f. The timestamp t and location l are straightforward representations of date/time and latitude/longitude coordinates, while features f can represent multiple aspects of different types of social media, e.g., text in a tweet, user tags for a photo, etc.

Definition 2 (Data Stream): Building upon Definition 1 (Social Media Post), we now have multiple social media posts arriving in a real-time data stream. Let $\mathcal{S} = \{p_1, p_2, \ldots, p_n\}$ denote the first n posts from the data stream, order temporally such that for p_i and p_j where $i < j$, $t_i \leq t_j$.

Definition 3 (Current/Query Window): In the context of a Data Stream $\mathcal{S} = \{p_1, p_2, \ldots, p_n\}$, we define a current/query window $\mathcal{W} = \{p_{n-m+1}, \ldots, p_{n-1}, p_n\}$, where $\mathcal{W} \subseteq \mathcal{S}$. This current/query window represents the current set of social media posts from post p_{n-m+1} to post p_n. For generalizability, the window size can be based on either a fixed number of posts, $m > 0$, or a fixed duration between posts p_{n-m+1} and post p_n, i.e. $t_n - t_{n-m+1}$.

B. Formal Problem Definition

The focus of our work is to develop an algorithm for detecting spatio-temporal events from streaming social media, based on a provided set of current social media posts, i.e., the query/current window. We define a spatio-temporal event as a set of social media posts that represents an increase in activity across a period of time within the same locality, based on the current/query window.

Given a data stream of social media posts $\mathcal{S} = \{p_1, p_2, \ldots, p_n\}$ and a query window $\mathcal{W} = \{p_{n-m+1}, p_{n-m+2}, \ldots, p_n\}$ that represents currently observed social media posts, we want to identify a set of posts $\mathcal{E} \in \mathcal{W}$ with the following goals:

- Spatial Proximity, e.g. $\sum_{p_x \in \mathcal{E}} \sum_{p_y \in \mathcal{E}} dist(l_x, l_y)$ should be significantly smaller than that for the same number of posts drawn uniformly at random from $\mathcal{W}$.
- Temporal Proximity, e.g. $\sum_{p_x, p_y \in \mathcal{E}} (t_y - t_x)$ (for consecutive p_x and p_y), should be significantly smaller than that for the same number of posts drawn uniformly at random from $\mathcal{W}$.
- Significance, $|\mathcal{E}|$ should be as large as possible while maintaining Spatial and Temporal Proximity goals.

In short, we are selecting a subset of social media posts that are representative of a spatio-temporal event, based on their spatial and temporal proximity.

IV. Proposed Algorithm

In our work we address spatial proximity by considering windows that are defined in terms of a region, and we address temporal proximity by considering a sliding window and assessing the change in the number of posts in a given region for two consecutive windows. Formally we consider a number of sliding windows, $\mathcal{W}_i(\Gamma_j)$, identified by an unbounded slide sequence number $i = 1, 2, \ldots$ and finite set of regions Γ_j, $j = 1, 2, \ldots, \gamma$, where $\bigcup \Gamma_j = \Gamma$. All of our sliding windows have a time duration of T and slide increment ΔT, with the head of the window being T_i, i.e. every sliding window i covers time interval $[T_i - T, T_i)$. In this way, we define

$$\mathcal{W}_i(\Gamma_j) = \mathcal{W}_{ij} = \{p_x \mid T_i - T \leq t_x < T_i, l_x \in \Gamma_j, p_x \in \mathcal{S}\}\,. \tag{1}$$

Each region serves as a spatial proximity bound for the posts that it contains, in the sense that we can consider the posts being within a given region as satisfying the spatial proximity goal from the problem definition. There are many ways that regions can be selected, e.g. they could be a uniform mesh based partition of the space, or each region could be associated with a point of interest (POI) in the space (e.g. a region around a park or building), etc. In our work we consider a multi-scale region selection approach based on a quad-tree division of space; in this case regions are overlapping with some regions subsuming others. In previous work we have also considered the POI region selection approach and we make comparisions between them in this paper.

In order to assess the change in the number of posts from one window to the next, we assume that the number of posts arriving in a given time interval has a Poisson distribution, and we thereby assign an estimate ΔT-arrival rate of posts for each region based on its sliding window:

$$\lambda_{ij} = \frac{|\mathcal{W}_{ij}|\Delta T}{T}, \tag{2}$$

where $|x|$ is the cardinality of set x. Finally, as the basis for event detection in each region, for each window $\mathcal{W}_{ij}$, we consider the observed number of posts in the slide increment interval $[T_i, T_i + \Delta T)$,

$$C_{ij} = \left| \left\{ p_x \mid T_i \leq t_x < T_i + \Delta T, l_x \in \Gamma_j, p_x \in \mathcal{S} \right\} \right|,$$

and make use of the Poisson p.m.f.:

$$\mathbb{P}\big[C_{ij}; \lambda_{ij}\big] = P_{ij} = e^{\lambda_{ij}} \frac{\lambda^{C_{ij}}}{C_{ij}!}. \tag{3}$$

If P_{ij} is significantly low (below a threshold) then we consider the possibility that region j has exhibited an event and we consider the posts within the slide increment to potentially be comprising that event. The details of our approach, called Spatio-temporal Online Event Detection Algorithm, includes more aspects that are explained next: *(1)* building a multiscale spatial resolution grid using the quad-tree method, *(2)* event detection using the Poisson model and signal smoothing, *(3)* event merging and *(4)* event pruning. Generally, our algorithm maintains an unbounded set of detected events $\mathcal{E}$ found in the unbounded stream $\mathcal{S}$. A detailed explanation for each phase is provided in the following subsections. As well, Table II provides an overview of the notation used in the algorithm.

A. Phase 1: Build Quad-tree

In this phase, we use the quad-tree method for spatial decomposition [28], [29]. It has been used in a variety of applications including image processing, computer graphics, geographic information systems and robotics [30], [31]. We construct

TABLE II
ADDITIONAL NOTATION

Γ_j	region j where $\Gamma = \bigcup \Gamma_j$ contains all posts in $\mathcal{S}$
$\mathcal{W}_{ij}$	the set of posts in the sliding window at interval i for region j
C_{ij}	number of posts in the window slide increment interval i for region j
λ_{ij}	estimate rate of posts at interval i for region j
P_{ij}	Poisson signal at interval i for region j
τ_1	Poisson signal threshold
F_{ij}	event signal at interval i for region j
τ_2	event detection threshold
α	event signal decay parameter
$\theta_{duration}$	event duration threshold
θ_{entity}	minimum entities threshold
θ_{area}	quad-tree node region area threshold
θ_{count}	quad-tree node post count threshold

a quad-tree at each time interval i. The quad-tree in two dimensional space starts with a large rectangular region, in our work $\Gamma_1 = \Gamma$, which represents the root of the quad-tree. The root region Γ_1 is subdivided into four equal sized regions $\{\Gamma_{11}, \Gamma_{12}, \Gamma_{13}, \Gamma_{14}\}$ and each subregion is recursively subdivided, i.e. creating $\{\Gamma_{111}, \Gamma_{112}, \dots\}$, and so on. Subdivision of a region x only occurs if both $|\mathcal{W}_{ix}| > \theta_{count}$ posts and the area of region x is at least θ_{area}. These constraints limit the minimum spatial resolution. As the quad-tree is constructed we also compute λ_{ij} and C_{ij} for each node, including internal nodes; here node is synonomous with region in that region j is node j.

B. Phase 2: Event Detection

For a sliding window interval i and all regions j (including those at internal nodes of the quad-tree), we use the Poisson distribution [32]–[34] to measure how likely the observed number of posts, C_{ij}, is for the slide increment ΔT that immediately follows the sliding window. The estimate arrival rate of posts is computed as in Eq. 2 and the probability, P_{ij}, of observing C_{ij} posts in time ΔT is computed as in Eq. 3. The more unlikely the observation, which may result from a significantly large increase or decrease in posts from the mean, the more we consider the posts (or lack thereof) to comprise an event. Therefore regions with $P_{ij} < \tau_1$, a constant threshold, could be flagged as potential regions for events. To compensate for sparse and/or incomplete data, where the stream of posts may not have a significantly strong representation of social media posts, we "smooth" the Poisson signal by computing an exponential decaying average *event signal*, F_{ij}:

$$\delta_{ij} = \begin{cases} \frac{\tau_1 - P_{ij}}{\tau_1}, & \text{if } P_{ij} < \tau_1 \\ 0, & \text{otherwise} \end{cases} \qquad (4)$$
$$F_{ij} = \alpha F_{i-1,j} + (1-\alpha)\delta_{ij}$$

where δ_{ij} is the scaled Poisson signal, F_{ij} and $F_{i-1,j}$ are the event signal values for node (region) j at the i and $i-1$ intervals respectively, and $0 \le \alpha \le 1$ is a constant decay parameter. Finally, if $F_{ij} \ge \tau_2$, an event detection threshold, we flag the posts, or more specifically the interval and region, as comprising an event.

C. Phase 3: Merge Events

Each event found in the previous Event Detection phase has a different spatial resolution and a fixed temporal scale (ΔT). In this phase, we construct events with multiscale temporal resolution using a merging method. Events at the same region that occurr at consecutive time intervals are merged. This gives an estimate for the period of time during in which an event is highlighted rather than assuming a predefined fixed duration (ΔT). For instance, if two events e_1 and e_2 occur in the same region at time intervals i and $i+1$ in order, then both events are combined to one event with period of $2\Delta T$. When merging we combine the posts and average the signal strength for the merged event.

D. Phase 4: Prune Events

To further increase the precision of our event detection and to handle spatio-temporal events that occur over a changing resolution we prune events after merging them. First, we only select events with duration $\ge \theta_{duration}$ to be included in the final set of detected events. The idea is that the longer the event duration, the more reliable and accurate it is. In other words, regions/nodes which are flagged for short periods are most likely to be noise (i.e. false positives). Second, the fact that we compute the signal for all quad-tree nodes (i.e. both internal and leaf nodes), leads to the propagation of some flagged events over the different tree levels (i.e. multi spatial resolution). So if an event is detected at the same time on different tree levels, we only keep the node with the strongest signal. In other words, if overlapping tree nodes (i.e. parent, child, grand child, and so on) are flagged as events from time t_1 to t_2, then we select the node (i.e. region) with the strongest signal to be the spatial resolution of the final detected event. This gives us a set of unique events which happened at different spatial and temporal resolutions. Third, we utilize the entities in the social media posts to detect and eliminate spam or fake events. We extract the set of unique entities (which may be keywords, mentions, hashtags, etc., depending on the type of social media post) across the posts in the event. If the size of the set is less than θ_{entity} then we remove the event.

V. Experimental Design

In this section, we describe our datasets and give an overview of the evaluation metrics and baseline algorithm used in our experimental methodology.

A. Dataset and Data Collection

To demonstrate the generalizability of our proposed algorithm, we perform our experimental evaluation on two datasets based on Twitter and Flickr. For our Twitter dataset, we performed a two-stage collection of tweets, similar to [35], [36]. We first used the Twitter REST API to retrieve all geo-tagged tweets posted by users in Melbourne in 2017. As we focus on geo-tagged tweets, this collection process resulted in 203519 geo-tagged tweets by 22264 different users. For our Flickr dataset, we focused on geo-tagged photos posted in Melbourne, based on the Yahoo! Flickr 100M Creative Commons (YFCC100M) dataset [37], [38]. The YFCC100M dataset comprises 100M geo-tagged photos and videos along with their meta-data such as latitude/longitude coordinates, date/time taken, photo name, user description, assigned tags, etc.

B. Evaluation Methodology

*1) **Precision**:* We use precision to measure the ratio of correctly detected events (true positives) to the total detected events. The absence of the ground truth makes the task of computing precision very difficult. As it is impractical to manually label the overly large number of events in the dataset, we propose a semi-automated assessment methodology using Google search results where each event is assigned 1 if it is true event, 0 otherwise. To do so, we first query Google using the top k entities as well as the date-time of each detected event. We use Google query results to decide whether an event is True or False, 1 and 0 refer to true and false event, respectively. If we don't get any useful information about the event from Google, then we manually look at the posts of the event to decide if it is a personal/private event, spam or wrong event.

*2) **Recall**:* Recall reflects the ability of the model to find all actual events within a dataset. In the context of event detection, recall measures the percentage of detected events with respect to important events/news appearing on a real-world news headlines. Similar to precision, we do a manual assessment for recall due to the absence of ground truth events. This is done by using Google search engine to select the most common events appearing on the news headlines for the days corresponding to the analysis. This includes festivals, public holiday events and international performances occur in the area of analysis. Each event is represented by a list of entities, which are used to manually decide whether an event is detected by our method or not.

*3) **Strength Index (SI)**:* To examine if the posts assigned to an event $e = (region, start, end, period, posts, signal)$ are relevant or not, we introduce a metric, which we refer to as the event strength index (SI). SI is the fraction of the retrieved top entities to the total count of event posts. We use SI as an indicator of how important/precise a reported event is. For an event e with total number of posts $C = |posts|$ and χ_i being the i-th most frequent entity (could be hashtags and mentions for twitter or image tags and description for Flickr), we calculate SI using the following formula:

$$\text{Strength Index (SI)} = \frac{\sum_{i=1}^{k} C_{\chi_i}}{C}, \tag{5}$$

for constant $k > 0$, where C_{χ_i} is the number of posts that contain χ_i. SI ranges from 0 to k, where k is the number of top entities. We obtain a small value for SI ($\ll 1$), when the top entities do not match the context of the detected event or if they are relevant but with a small number of occurrences. For example, a value of 0 for SI means that all posts for an event e are irrelevant, while a value of k means that all event posts contain at least one occurrence of each top entity.

C. Baseline Algorithm

To show the effectiveness of the proposed method, we compare it with a baseline algorithm that uses Points of Interests (POIs). POIs have been frequently used in location-based recommendation [39]–[41] and we develop a baseline using the similar idea of tagging geo-tagged social media to POIs. Similar to earlier works, we obtain a list of known and popular POIs for each city from their respective Wikipedia entries. In this baseline algorithm, we utilize a spatial representation of tweets based on their proximity (<100m) to known POIs, instead of assigning tweets to dynamically-sized grids based on quad trees. The remaining steps of computing Poisson signals and determining event duration remain the same as previously described in Section IV.

1) Incremental Clustering for Real-time Event Detection: Among the existing event detection techniques and algorithms discussed in Section II, we select the event detection approach proposed in [7] for comparison with our approach. The reason is that it is very closely related to our introduced problem. The approach detects significant clusters that are sufficiently dense and large, in streams of spatial events with the advantage of tracking cluster evolution over time.

Given a list of active data points (i.e. spatial events) that occurred in the interval $[t_c - \Delta T, t_c]$, where t_c is the current time and ΔT is a maximal temporal gap, the algorithm finds the set of significant clusters by repeatedly extracting a set of event circles and unions every time tick (t). An event circle C is a group of active events that fits in a circle with maximal radius R. While a union is a set of event circles that have at least k-overlapped events (i.e. k-connecting events). Finally, a significant event cluster is a union that includes at least N spatial events, i.e. minimal cluster size. The values for parameters t, ΔT, R, K, and N are user-specified. More details about the algorithm can be found in [7].

We implemented the full algorithm in Python and experimented using the geotagged tweets dataset for Melbourne in 2017.

TABLE III
PARAMETERS USED IN THE PROPOSED METHOD

Method	***Parameter = Value***
Quad-tree	θ_{count}= 20
	θ_{area} = 0.001 sqkm
Poisson model	T = 3 days
	ΔT = 10 minutes
	signal threshold τ_1 =0.01
	event detection threshold τ_2 =0.4
	event signal decay coefficient α = 0.5
Event filtering	duration threshold $\theta_{duration}$= 50 minutes
	minimum top entities θ_{entity} = 2

TABLE IV
EVENT DETECTION RESULTS FOR FLICKR AND TWITTER

Case Study	***Twitter***	***Flickr***
Month-Year	2017	Jan-2013
Number of users	22264	90
Number of posts:	203519	995
Total 10-min intervals & nodes	15253260	230468
Total flagged nodes	29520	726
Merging adjacent nodes	25108	165
Duration filtering	299	19
Filtering propagated events	158	7
Filtering spam events	137	7

VI. RESULTS

In this section, we evaluate the proposed method in four different aspects. Firstly, we present a preliminary analysis of the proposed method (Section VI-A). Secondly, we present a detailed comparative analysis with the baseline algorithm (Section VI-D1). Thirdly, using the tweets over a period of one-year we evaluate our algorithm based on the precision, recall and strength index as statistical metrics (VI-B). Finally, we show a case study of event detection using Flickr image dataset (Section VI-C).

A. Preliminary Analysis

In this section, we used a subset of the collected tweets to evaluate the individual phases of the proposed method. We extracted January-2017 Melbourne tweets which contains 23327 geotagged tweets for 5427 users. First, quad-tree is used to construct multiscale spatial grid. Then, events are detected using Poisson model. Following this, a smoothing function is applied for accurate estimation of event duration. Finally, a false positives removal phase is performed to eliminate both falsely highlighted events and spam events. Table III presents the different parameters used in the proposed method. The parameters are chosen after several experiments, to achieve the best performance.

B. Case Study: Twitter Dataset

To evaluate the performance and reliability of the proposed method, we experiment the whole dataset for Melbourne in 2017. Figure 1 visualises some of the detected events on the map. Each event has the start and end time, total tweets, area and the top 5 hashtags/mentions. We use the top hashtags/mentions along with the event time to manually evaluate the correctness of the event. Table IV, column ”Twitter” shows the total number of flagged events after each phase of the proposed method. In total, we detect 137 events after the removal of all false positives.

1) ***Precision and Strength Index Results****:* We select a random 45 events as the evaluation set. Table V reports the manual evaluation results for sample selected events from the evaluation set. For each event, the algorithm results the event start and end date-time, area/region in $sqkm$, tweets count, top 5 hashtags/mentions with its occurrences. In total, 40 out of 45 are correct events according to the manual evaluation, with a precision measure of 89%. False positives are highlighted in red in Table V. The table shows that both local and global events are detected using the proposed method (see the ”area” column in Table V).

SI index is also reported for all events (see column ”SI” in Table V). We obtain SI with an average of 1.2 across the evaluation set. This is an indication that all tweets for an event contain at least one of the relevant hashtags/mentions, which confirms the accurate results of the proposed method.

2) ***Recall and Strength Index Results****:* We select 15 events to assess the recall with a total of 10 correctly detected events according to the manual evaluation, with a recall measure of 66.7% and average SI of 1.034. The reduction in recall is explainable since the social media does not contain information about all actual events. This causes a certain increase in the false negatives. Table VI reports the date-time, top 5 entities, tweets count, event area and SI for sample selected events from the evaluation set. In the table, the false negatives are highlighted in blue.

C. Case Study: Flickr Dataset

In this section we present our second case study for event detection using the Flickr dataset introduced in Section V-A. The dataset was further reduced by keeping only geotagged images for January-2013 for Melbourne. The chosen years have largest number of photos taken in January. We evaluated the proposed event detection method using the set of images collected for each city. In our experiments, instead of extracting top k hashtags/mentions as in the twitter case study, we use the title, userTags and description attributes for each image to extract the most frequent entities of an event. Table IV, column ”Flickr” reports the results each individual

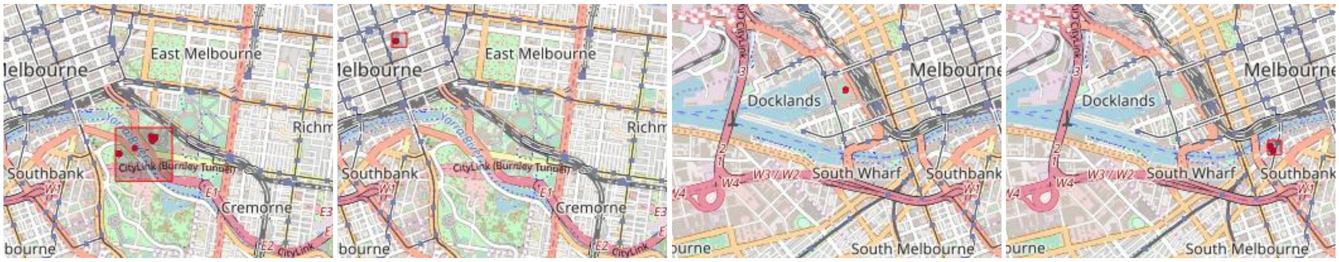

Start: 01-28 09:40	End: 01-28 10:40	Start: 02-27 00:00	End: 02-27 01:00	Start: 03-10 13:50	End: 03-10 14:40	Start: 04-23 06:00	End: 04-23 06:50
Tweets count: 35	Area (sqkm): 0.44	Tweets count: 17	Area (sqkm): 0.028	Tweets count: 13	Area (sqkm): 0.002	Tweets count: 10	Area (sqkm): 0.028
Top 5:	#nickcave: 17.1%	**Top 5:**	#greekandthecity: 82.4%	**Top 5:**	#purposetour: 46.2%	**Top 5:**	#tvweeklogies: 50 %
#australianopen: 11.4%	@serenawilliams: 8.6%	#lovelonsdale: 76.5%	#kkfuntimes: 76.5%	#justinbieber: 30.8%	#melbourne: 23.1%	#redcarpet: 20%	@nazeem: 10%
#sidneymyermusicbowl: 8.6%	@australianopen: 8.6%	#kkrockchic: 64.7%	#mc: 17.6%	#purposeworldtour: 15.4%	#belieber: 15.4%	@rebeccavallance: 10%	@hishandsomeself: 10%

Fig. 1. Sample detected events using Twitter Data in Melbourne, 2017

TABLE V
PRECISION RESULTS: MELBOURNE DETECTED EVENTS IN 2017 USING GEOTAGGED TWEETS (DUE TO SPACE CONSTRAINTS, 5 EVENTS ARE SHOWN)

Detected Events Results						**Manual Assessment**	
ID	**Date & Time**	**Top 5 hashtags/mentions: # of tweets**	**Count**	**Area**	**SI**	**True?**	**Event Description**
5	-S 01-28 09:40-E 01-28 10:40	#nickcave: 6, #australianopen: 4, #23: 4, @serenawilliams: 3, @australianopen: 3	35	0.44	0.57	1	Serena Williams wins Australian Open final (Tennis)
9	-S 02-27 00:10-E 02-27 01:00	#greekandthecity: 13, #lovelonsdale: 13, #kkfuntimes: 13, #kkrockchic: 11, #mc: 3	18	0.44	2.94	1	Thousands turn out to celebrate Lonsdale St Festival's 30th anniversary
12	-S 03-18 06:50-E 03-18 07:40	#adele: 10, #melbourne: 4, #etihadstadium: 3, #adele2017: 2, #benandjase: 1	14	0.44	1.43	1	Adele adds new Melbourne show to 2017 Australian tour
22	-S 04-05 10:00-E 04-05 12:20	#worlds50best: 41, @theworlds50best: 28, @australia: 14, #seeaustralia: 11, #restaurantaustralia: 10	94	0.03	1.11	1	The 2017 awards ceremony for the World's 50 Best Restaurants 2017
26	-S 04-25 04:00-E 04-25 05:40	#anzacday: 15, #lestweforget: 8, @mcg: 8, #mcg: 6, #afldonspies: 5	44	0.01	0.95	1	* Melbourne Cricket Ground: AFL -round 5 - Essendon VS Collingwood. * ANZAC Day public holiday

phase of the proposed method using Flick Dataset for each city.

D. Comparative Analysis of Baseline Algorithms

1) POI-based Event Detection: We use January-2017 Melbourne geotagged tweets dataset to run this experiment. We selected 242 POIs (100m x 100m) in Melbourne where most of events occur. We apply the proposed event detection algorithm on the POIs instead of the the multiscale grid generated by quad-tree method. Every ΔT (i.e. 10-min), we flag all POIs with smoothed signal less than the threshold (τ_2).

Table VII reports the details of the detected events on $28/01/2017$ using each of the proposed and baseline methods. The results show that our approach is able to detect all events as identified by the baseline method, along with additional events with different spatial scales that the baseline method was unable to detect. Also, the conducted experiments show that POI based method detects events with shorter duration than the quad-tree based method. The reason is that POI grid cells are is small (100m x 100m).

2) Incremental Clustering for Event Detection: We run the algorithm using the following parameters: time tick $t = 10$-min, $R = 0.25km$, $\Delta T = 3$hrs, $K = 2$, and $N = 5$. The values of the first two parameters are chosen to be close to the values used in our approach for fair comparison. While the values of the other parameters are chosen to be the same as the one used in the paper [7]. For each time tick, i.e. 10min, we record the set of significant unions (i.e. clusters). In total, we obtained 154546 significant events in Melbourne 2017. To avoid detecting same event more than once, we group the set of detected events by their union ID. Then we keep only the event with the highest number of tweets for each union ID. This grouping strategy decreased the total number of significant events to 3994, since it removes all redundant events.

For quantitative comparison with our approach, we compute precision and recall for the implemented algorithm. For precision, we select a random 20 events as the evaluation set. For each event, we calculate the the event start and end date-time by using date-time of all tweets assigned to each cluster. We also report area/region in $sqkm$, tweets count, top 5 hashtags/mentions with its occurrences, description and SI index for all events. In total, 9 out of 20 are correct events according to the manual evaluation, with a precision measure of 45%. It is clear that the precision is very low compared to our results. The reason is that using small values for k and N resulted in large number of detected events, where the

TABLE VI
RECALL RESULTS: MELBOURNE COMMON EVENTS OCCURRED IN 2017

#	Event Description	Date-time	Entity: occurrences	*	Area	SI
1	The Night Market; Wed 6-9pm	-S 01-18 10:10 -E 01-18 11:00	#melbourne: 3, #queenvictoriamarket: 2, #glutenfree: 1, #nofilter: 2, @hunde: 1	35	28.181	0.258
2	Moomba Festival; 10-13 Mar	-S 03-12 01:30 -E 03-12 01:50	#melbourne: 4, #davidhockney: 1, #shrineofremembrancemelbourne: 1, #ladiesinblack: 1, @moombafestival: 1	16	7.04	0.498
5	Patti Smith performs; Apr 16	-S 04-16 13:20 -E 04-16 14:00	#music: 4, #livemusic: 4, #melbourne: 4, #horses: 3, #pattismith: 3	8	0.028	2.25
6	Australia Day; Jan 26	-S 01-26 05:40 -E 01-26 06:10	#art: 6, @ngv: 5, #ipad: 5, @australianopen: 5, #inspiration: 5	20	1.76	1.3
7	Labour Day; Mar 13	*Not Detected*				
9	Easter Sunday; Apr 16	-S 04-17 04:00 -E 04-17 04:20	#brunch: 1, #eastersunday: 1	4	7.04	0.5
10	ANZAC Day; Apr 25	-S 04-25 04:00 -E 04-25 05:40	#anzacday: 15, #lestweforget: 8, @mcg: 8, #mcg: 6, #afldonspies: 5	44	0.007	0.96
11	Queen's Birthday; Jun 12	*Not Detected*				
13	Melbourne Cup; Nov 7	-S 11-07 06:40-E 11-07 07:10	#melbournecup: 5, #foodlover: 1, #burger: 1, #foodie: 1, #melbourne: 5	11	1.76	1.183
15	Boxing Day; Dec 26	-S 12-26 02:00-E 12-26 02:40	#boxingdaytest: 6, #beatengland: 5, #mcg: 4, #ashes2017: 3, #kneipping: 3	10	0.007	2.1

TABLE VII
MELBOURNE DETECTED EVENTS IN 28-JANUARY 2017 USING EACH OF POI AND QUAD-TREE METHODS

#	Time - Duration	**Top 5 hashtags/mentions** (hashtag/mention : % of occurrences)	Count	Area
		POI events		
1	-S 08:30 -D 40	#ausopen: 37.5, @australianopen: 25, #williamssisters: 18.8, @serenawilliams: 12.5, #melbourne: 12.5	16	0.01
2	-S 10:00 -D 20	#nickcave: 80, #sidneymyermusicbowl: 60, #livemusic: 20, #melbourne: 20	5	0.01
3	-S 11:10 -D 40	#ausopen: 50, #womensfinal: 33.3, #rodlaverarena: 33.3, #ausralianopen: 25, #serena: 25	12	0.01
		Quad-tree events		
1	-S 08:30 -D 40	#ausopen: 33.3, #williamssisters: 23.8, @australianopen: 23.8, @venuseswilliams: 14.3, @serenawilliams: 14.3	21	0.052
2	-S 08:40 -D 30	#ausopen: 50, #williamssisters: 25, @australianopen: 25, #melbourne: 16.7, #grandslam: 8.3	12	0.003
3	-S 10:00 -D 20	#nickcave: 80, #sidneymyermusicbowl: 60, #livemusic: 20, #melbourne: 20	5	0.013
4	-S 10:00 -D 40	#nickcave: 20.8, @serenawilliams: 12.5, #23: 16.7, #australianopen: 12.5, #sidneymyermusicbowl: 12.5	24	0.834
5	-S 10:00 -D 50	#nickcave: 13.2, @serenawilliams: 10.5, #sidneymyermusicbowl: 7.9, #23: 10.5, #australianopen: 7.9	38	13.347
6	-S 10:20 -D 20	@australianopen: 28.6, #serenavsvenus: 14.3, #23: 28.6, #venus: 14.3, #williamssisters: 14.3	7	0.003
7	-S 11:10 -D 30	#ausopen: 54.5, #womensfinal: 36.4, #rodlaverarena: 36.4, #ausralianopen: 27.3, #serena: 27.3	11	0.003
8	-S 20:50 -D 20	#melbourne: 100, #hiring: 53.8, #job: 53.8, #bourkestreet: 46.2, #careerarc: 46.2	13	0.052

TABLE VIII
RECALL RESULTS FOR CLUSTERING BASED EVENT DETECTION: MELBOURNE COMMON EVENTS OCCURRED IN 2017 (DUE TO SPACE CONSTRAINTS, 5 EVENTS ARE SHOWN)

#	Event Description	Date-time	Entity: occurrences	*	Area	SI
1	The Night Market; Wed 6-9pm	-S 08-23 8:40 -E 08-23 10:20	#melbourne: 17, #moomba: 6, #vscocam: 3, #vsco: 3, #job: 3	65	0.25	0.5
2	Moomba Festival; 10-13 Mar	-S 03-12 1:00 -E 03-12 3:50	#melbourne: 4, #davidhockney: 1, #shrineofremembrancemelbourne: 1, #ladiesinblack: 1, @moombafestival: 1	16	7.04	0.498
3	Melbourne Food and Wine Festival; Mar 31 – Apr 9	*Not Detected*				
4	White Night Melbourne; Feb 18	-S 02-18 10:30 -E 02-18 13:20	#whitenightmelb: 3, #whitenight: 5, #victoria: 1, #melbournelife: 1, 0	7	0.25	1.43
5	Patti Smith performs; Apr 16	-S 04-16 10:50 -E 04-16 13:40	#melbourne: 5, #music: 4, #pattismith: 4, #livemusic: 4, #horses: 3	12	0.25	1.67

majority of these events are just noise. This can be improved by increasing the values of k, and N which will result in much less number of detected events and accordingly less false positives.

For recall, we use the same 15 common events used for recall assessment in our approach. A total of 13 correctly detected events according to the manual evaluation, with a recall measure of 86.7% which is higher than the recall measure of our approach. The reason is that the parameters used in this experiment results in large number of true negatives as well as large number of false positives. Parameters should be tuned to balance the trade off between recall and precision. Table VIII reports the date-time, top 5 hashtags/mentions with its occurrences, tweets count, event area and SI for all events. The false negatives are highlighted in blue in the table.

VII. CONCLUSION

In this paper, we present a multiscale spatio-temporal real-time event detection approach which is capable of detecting social media events of different spatial and temporal resolution in real-time. Also, the proposed method does not require a list of defined topics for event detection and effectively detects both local and global events. The method is evaluated using two different social media datasets: Twitter and Flickr. The experiments have demonstrated that the proposed method achieves better results than the baseline algorithm. In the future, we plan to improve our method by taking into account the changing structure of the constructed quad-tree over time. Also, more experiments will be conducted to fine-tune the parameters of the proposed method using different datasets. The proposed method will be extended to use non-geotagged social media data based on textual information [42]–[46]. Finally, we can also improve tour recommendation works by planning itineraries that avoid detected events such as accidents [47]–[51].

ACKNOWLEDGEMENTS

This research is funded in part by the Defence Science and Technology Group, Edinburgh, South Australia, under contract MyIP:7293, and the Singapore University of Technology and Design under grant SRG-ISTD-2018-140.

REFERENCES

[1] C. C. Aggarwal and K. Subbian, *Event Detection in Social Streams*, pp. 624–635. [Online]. Available: https://epubs.siam.org/doi/abs/10.1137/1.9781611972825.54

[2] R. Popovici, A. Weiler, and M. Grossniklaus, "On-line clustering for real-time topic detection in social media streaming data," in *SNOW 2014 Data Challenge*, 2014, pp. 57–63.

[3] C. Zhang, G. Zhou, Q. Yuan, H. Zhuang, Y. Zheng, L. Kaplan, S. Wang, and J. Han, "Geoburst: Real-time local event detection in geo-tagged tweet streams," in *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*. ACM, 2016, pp. 513–522.

[4] M. Walther and M. Kaisser, "Geo-spatial event detection in the twitter stream," in *European conference on information retrieval*. Springer, 2013, pp. 356–367.

[5] X. Dong, D. Mavroeidis, F. Calabrese, and P. Frossard, "Multiscale event detection in social media," *Data Mining and Knowledge Discovery*, vol. 29, no. 5, pp. 1374–1405, 2015.

[6] J. Capdevila, G. Pericacho, J. Torres, and J. Cerquides, "Scaling dbscan-like algorithms for event detection systems in twitter," in *Algorithms and Architectures for Parallel Processing - 16th International Conference, ICA3PP 2016, Granada, Spain, December 14-16, 2016, Proceedings*, 2016, pp. 356–373. [Online]. Available: https://doi.org/10.1007/978-3-319-49583-5_27

[7] N. Andrienko, G. Andrienko, G. Fuchs, S. Rinzivillo, and H. Betz, "Detection, tracking, and visualization of spatial event clusters for real time monitoring," in *2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, Oct 2015, pp. 1–10.

[8] X. Wang, F. Zhu, J. Jiang, and S. Li, "Real time event detection in twitter," in *International Conference on Web-Age Information Management*. Springer, 2013, pp. 502–513.

[9] M. Hasan, M. A. Orgun, and R. Schwitter, "Real-time event detection from the twitter data stream using the twitternews+ framework," *Information Processing & Management*, 2018.

[10] C. Zhang, D. Lei, Q. Yuan, H. Zhuang, L. Kaplan, S. Wang, and J. Han, "Geoburst+: effective and real-time local event detection in geo-tagged tweet streams," *ACM Transactions on Intelligent Systems and Technology*, vol. 9, no. 3, p. 34, 2018.

[11] Y. Huang, Y. Li, and J. Shan, "Spatial-temporal event detection from geo-tagged tweets," *ISPRS International Journal of Geo-Information*, vol. 7, no. 4, p. 150, 2018.

[12] K. H. Lim, S. Karunasekera, A. Harwood, and L. Falzon, "Spatial-based Topic Modelling using Wikidata Knowledge Base," in *Proceedings of the 2017 IEEE International Conference on Big Data (BigData'17)*, Dec 2017, pp. 4786–4788.

[13] C. Zhang, L. Liu, D. Lei, Q. Yuan, H. Zhuang, T. Hanratty, , and J. Han, "Triovecevent: Embedding-based online local event detection in geo-tagged tweet streams," in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2017, pp. 595–604.

[14] Y. Li, Q. Li, and J. Shan, "Discover patterns and mobility of twitter users—a study of four us college cities," *ISPRS International Journal of Geo-Information*, vol. 6, no. 2, p. 42, 2017.

[15] W. Qardaji, W. Yang, and N. Li, "Differentially private grids for geospatial data," in *2013 IEEE 29th International Conference on Data Engineering (ICDE)*, April 2013, pp. 757–768.

[16] A. Cui, M. Zhang, Y. Liu, S. Ma, and K. Zhang, "Discover breaking events with popular hashtags in twitter," in *Proceedings of the 21st ACM international conference on Information and knowledge management (CIKM'12)*, 2012, pp. 1794–1798.

[17] A. Zubiaga, D. Spina, R. Martinez, and V. Fresno, "Real-time classification of twitter trends," *Journal of the Association for Information Science and Technology*, vol. 66, no. 3, pp. 462–473, 2015.

[18] W. Xie, F. Zhu, J. Jiang, E.-P. Lim, and K. Wang, "Topicsketch: Real-time bursty topic detection from twitter," *IEEE Transactions on Knowledge and Data Engineering*, vol. 28, no. 8, pp. 2216–2229, 2016.

[19] A.-M. Popescu and M. Pennacchiotti, "Detecting controversial events from twitter," in *Proceedings of the 19th ACM international conference on Information and knowledge management (CIKM'10)*, 2010, pp. 1873–1876.

[20] S. Dori-Hacohen and J. Allan, "Detecting controversy on the web," in *Proceedings of the 22nd ACM international conference on Conference on information and knowledge management*, 2013, pp. 1845–1848.

[21] J. Weng and B.-S. Lee, "Event detection in twitter," in *Proceedings of ICWSM'11*, 2011, pp. 401–408.

[22] T. Sakaki, M. Okazaki, and Y. Matsuo, "Earthquake shakes Twitter users: real-time event detection by social sensors," in *Proceedings of the 19th international conference on World Wide Web*, 2010, pp. 851–860.

[23] ——, "Tweet analysis for real-time event detection and earthquake reporting system development," *IEEE Transactions on Knowledge and Data Engineering*, vol. 25, no. 4, pp. 919–931, 2013.

[24] R. Li, K. H. Lei, R. Khadiwala, and K. C.-C. Chang, "TEDAS: A twitter-based event detection and analysis system," in *Proceedings of the 28th International Conference on Data Engineering (ICDE'12)*, 2012, pp. 1273–1276.

[25] K. H. Lim, S. Jayasekara, S. Karunasekera, A. Harwood, L. Falzon, J. Dunn, and G. Burgess, "RAPID: Real-time Analytics Platform for Interactive Data Mining," in *Proceedings of the 2018 European Conference on Machine Learning and Knowledge Discovery in Databases (ECML-PKDD'18)*, Sep 2018.

[26] C. Li, A. Sun, and A. Datta, "Twevent: segment-based event detection from tweets," in *Proceedings of the 21st ACM International Conference on Information and Knowledge Management*, 2012, pp. 155–164.

[27] F. Atefeh and W. Khreich, "A survey of techniques for event detection in twitter," *Computational Intelligence*, vol. 31, no. 1, pp. 132–164, 2015.

[28] R. A. Finkel and J. L. Bentley, "Quad trees a data structure for retrieval on composite keys," *Acta informatica*, vol. 4, no. 1, pp. 1–9, 1974.

[29] S. Wang and M. P. Armstrong, "A quadtree approach to domain decomposition for spatial interpolation in grid computing environments," *Parallel Computing*, vol. 29, no. 10, pp. 1481–1504, 2003.

[30] H. Samet, "The quadtree and related hierarchical data structures," *ACM Computing Surveys (CSUR)*, vol. 16, no. 2, pp. 187–260, 1984.

[31] J. B. Rosenberg, "Geographical data structures compared: A study of data structures supporting region queries," *IEEE transactions on computer-aided design of integrated circuits and systems*, vol. 4, no. 1, pp. 53–67, 1985.

[32] T. Yamane, "Statistics: An introductory analysis," 1973.

[33] R. R. Sokal, F. J. Rohlf *et al.*, *The principles and practice of statistics in biological research*. WH Freeman and company San Francisco:, 1969.

[34] J. K. Patel, C. Kapadia, and D. B. Owen, "Handbook of statistical distributions," Tech. Rep.

[35] K. H. Lim, K. E. Lee, D. Kendal, L. Rashidi, E. Naghizade, S. Winter, and M. Vasardani, "The Grass is Greener on the Other Side: Understanding the Effects of Green Spaces on Twitter User Sentiments," in *Proceedings of the 2018 Web Conference Companion (WWW'18)*, Apr 2018, pp. 275–282.

[36] K. H. Lim, K. E. Lee, D. Kendal, L. Rashidi, E. Naghizade, Y. Feng, and J. Wang, "Understanding Sentiments and Activities in Green Spaces using a Social Data-driven Approach," in *Smart Cities: Issues and Challenges*. Elsevier, Jun 2019, pp. 77–107.

[37] Yahoo! Webscope, "Yahoo! Flickr Creative Commons 100M dataset (YFCC-100M)," 2014, http://webscope.sandbox.yahoo.com/catalog.php?datatype =i&did=67.

[38] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L.-J. Li, "YFCC100M: The new data in multimedia research," *Communications of the ACM*, vol. 59, no. 2, pp. 64–73, 2016.

[39] G. Cai, K. Lee, and I. Lee, "Itinerary recommender system with semantic trajectory pattern mining from geo-tagged photos," *Expert Systems with Applications*, vol. 94, pp. 32–40, 2018.

[40] D. Chen, C. S. Ong, and L. Xie, "Learning points and routes to recommend trajectories," in *Proc. of CIKM'16*, 2016, pp. 2227–2232.

[41] I. R. Brilhante, J. A. Macedo, F. M. Nardini, R. Perego, and C. Renso, "On planning sightseeing tours with TripBuilder," *Information Processing & Management*, vol. 51, no. 2, pp. 1–15, 2015.

[42] B. Han, P. Cook, and T. Baldwin, "Text-based twitter user geolocation prediction," *Journal of Artificial Intelligence Research*, vol. 49, 2014.

[43] L. Chi, K. H. Lim, N. Alam, and C. J. Butler, "Geolocation Prediction in Twitter Using Location Indicative Words and Textual Features," in *Proceedings of the 2nd Workshop on Noisy User-generated Text (WNUT'16)*, Dec 2016, pp. 227–234.

[44] K. H. Lim, S. Karunasekera, A. Harwood, and Y. George, "Geotagging Tweets to Landmarks using Convolutional Neural Networks with Text and Posting Time," in *Proceedings of the 24th International Conference on Intelligent User Interfaces Companion (IUI'19)*, 2019, pp. 61–62.

[45] A. Rahimi, T. Cohn, and T. Baldwin, "A Neural Model for User Geolocation and Lexical Dialectology," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL'17)*, 2017, pp. 209–216.

[46] W.-H. Chong and E.-P. Lim, "Exploiting user and venue characteristics for fine-grained tweet geolocation," *ACM Transactions on Information Systems*, vol. 36, no. 3, pp. 26:1–26:34, 2018.

[47] C. Chen, D. Zhang, B. Guo, X. Ma, G. Pan, and Z. Wu, "TripPlanner: Personalized trip planning leveraging heterogeneous crowdsourced digital footprints," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 3, pp. 1259–1273, 2015.

[48] Z. Friggstad, S. Gollapudi, K. Kollias, T. Sarlos, C. Swamy, and A. Tomkins, "Orienteering algorithms for generating travel itineraries," in *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. ACM, 2018, pp. 180–188.

[49] P. Padia, K. H. Lim, J. Chan, and A. Harwood, "Sentiment-Aware and Personalized Tour Recommendation," in *Proceedings of the 2019 IEEE International Conference on Big Data (BigData'19)*, Dec 2019.

[50] T. Liebig, N. Piatkowski, C. Bockermann, and K. Morik, "Dynamic route planning with real-time traffic predictions," *Information Systems*, vol. 64, pp. 258–265, 2017.

[51] ——, "Predictive trip planning-smart routing in smart cities," in *Proceedings of EDBT/ICDT Workshop on Mining Urban Data (MUD'14)*, 2014, pp. 331–338.