

Inferring Unusual Crowd Events from Mobile Phone Call Detail Records

Yuxiao Dong¹, Fabio Pinelli², Yiannis Gkoufas², Zubair Nabi²,
Francesco Calabrese² (✉), and Nitesh V. Chawla¹

¹ Department of Computer Science and Engineering,
University of Notre Dame, Notre Dame, USA
{ydong1,nchawla}@nd.edu

² IBM Research, Mulhuddart, Ireland
{fabiopin,yiannisg,zubairn,fcalabre}@ie.ibm.com

Abstract. The pervasiveness and availability of mobile phone data offer the opportunity of discovering usable knowledge about crowd behavior in urban environments. Cities can leverage such knowledge to provide better services (e.g., public transport planning, optimized resource allocation) and safer environment. Call Detail Record (CDR) data represents a practical data source to detect and monitor unusual events considering the high level of mobile phone penetration, compared with GPS equipped and open devices. In this paper, we propose a methodology that is able to detect unusual events from CDR data, which typically has low accuracy in terms of space and time resolution. Moreover, we introduce a concept of unusual event that involves a large amount of people who expose an unusual mobility behavior. Our careful consideration of the issues that come from coarse-grained CDR data ultimately leads to a completely general framework that can detect unusual crowd events from CDR data effectively and efficiently. Through extensive experiments on real-world CDR data for a large city in Africa, we demonstrate that our method can detect unusual events with 16% higher recall and over 10× higher precision, compared to state-of-the-art methods. We implement a visual analytics prototype system to help end users analyze detected unusual crowd events to best suit different application scenarios. To the best of our knowledge, this is the first work on the detection of unusual events from CDR data with considerations of its temporal and spatial sparseness and distinction between user unusual activities and daily routines.

1 Introduction

The ubiquity of mobile devices offers an unprecedented opportunity to analyze the trajectories of movement objects in an urban environment, which can

We wish to thank the Orange D4D Challenge (<http://www.d4d.orange.com>) organizers for releasing the data we used for testing our algorithms. Research was sponsored in part by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053 and the U.S. Air Force Office of Scientific Research (AFOSR) and the Defense Advanced Research Projects Agency (DARPA) grant #FA9550-12-1-0405.

have a significant effect on city planning, crowd management, and emergency response [4]. The big data generated from mobile devices, thus, provides a new powerful social microscope, which may help us to understand human mobility and discover the hidden principles that characterize the trajectories defining human movement patterns. Cities can leverage the results of the analytics to better provide and plan services for citizens as well as to improve their safety. For example, during the occurrence of expected or chaotic events such as riots, parades, big sport events, concerts, the city should be able to provide a proactive response in allocating the correct amount of resources, adapt public transport services, and more generally adopt all possible actions to safely handle such events. Many methods have been proposed in the literature to detect groups of people moving together from a trajectory database [14, 16, 18, 32], specifically the GPS data. However, only a very little percentage of people currently carry GPS devices, and share their movement trajectories with a central entity that can use them to identify crowd events.

In this paper, we study the problem of unusual event detection from mobile phone data that is opportunistically collected by telecommunication operators, in particular the Call Detail Records (CDR). In 2013, the number of mobile-phone subscriptions reached 6.8 billion, corresponding to a global penetration of 96%. The pervasiveness of mobile phones is spreading fast, with the number of subscriptions reaching 7.3 billion by 2014, from a recent report by International Telecommunications Union (ITU) at 2013 Mobile World Congress [7]. Therefore, CDR data represents a practical data source to detect and monitor unusual events considering the high level of mobile phone penetration. This is specifically useful in developing countries where other methodologies to gather crowd movement data (e.g., GPS or cameras) are very expensive to be installed.

The task of detecting unusual events from CDR data is very different from previous work on fine-grained trajectory data, such as GPS data, and presents several unique challenges. **Temporal sparseness:** CDR data only records the user location when a call or text message is made or received, thus is temporally sparse since call or message frequency of users is usually low and unpredictable. **Spatial sparseness:** The location information of users when they make a call or message is recorded as the location of the antenna, which brings the spatial sparseness of CDR data. **Non-routine events:** Our objective is to detect unusual crowd events from human daily movements, which mostly consist of usual routines. Thus, it is necessary to discriminate unusual crowd movements from routine trajectories.

To address these challenges, we aim to estimate the location of users in absence of spatio-temporal observations (i.e., the users don't make phone calls), detect groups of people moving together, and proactively discover unusual events. We propose a general framework to infer unusual crowd events from mobile phone data. Specifically, our contributions can be summarized as follows:

- We first define the cylindrical cluster to capture sparse spatio-temporal location data and provide practical methods to extract crowd events from CDR data, and further formalize the unusual crowd event detection problem by




Fig. 1. Process flow of the system to detect unusual events.

considering the similarity between individuals' trajectories and their historical mobility profiles.

- We provide a Visual Analytics Prototype System to help the end user (e.g. a city manager or analyst) analyze the detected crowd events and set the values of the parameters to best suit an application scenario.
- Finally, we evaluate our proposed framework on a real-world CDR dataset and demonstrate its effectiveness and efficiency. Our method significantly outperforms ($10\times$ precision and $+16\%$ recall) previous event detection methods on GPS data with verification on real-world unusual crowd events.

The mobile phone CDR data used in this work is collected from Cote d'Ivoire over five months, from December 2011 to April 2012. During that period, this Africa country faced the Second Ivorian Civil War and political crisis¹. From the election of new president and parliament, continued outbreaks of post-election conflicts happened, including boycott, violence and protest etc. The experimental results on this real-world dataset deliver the effectiveness of our proposed methodologies, which demonstrates the significant importance of our work in the supervision of unusual crowds and events for city and country management. Moreover, through the proposed method, city managers and officials can gain insights into non-ticketed events taking place in public spaces, which could lead to estimating the number of attendees and to estimating the event's success. A particular instance of such method has been recently implemented to help the police and the event organizers monitor visitors to the Mons 2015 - European Capital of Culture Opening Ceremony [3].

2 Unusual Event Detection Problem

Given the nature of CDR data, we face three major challenges in extracting accurate individual trajectories. First, we can only record user locations when they make calls or receive calls (text messages). As most mobile users do not make phone calls frequently and periodically, positions are not regularly sampled, as opposed to GPS navigation systems. Moreover, mobile users do not follow a call pattern consistently with others in the group. Second, when a user makes a call, CDR data only records the base station she is using, providing very low quality location information. Finally, the scenario that we are considering—e.g.,

¹ http://en.wikipedia.org/wiki/2010-11_Ivorian_crisis

going to a protest—is not consistent with an individual’s daily activity pattern such as going from home to office, thus we cannot leverage the previous history of the user to enrich his trajectory to make it more accurate.

We formally define the problem of unusual event detection and decompose the problem in different steps that enable us to solve the challenges brought by CDR data. Figure 1 shows the process flow to detect unusual events. The system receives CDR data as input, extracts clusters, and detects crowds from the sequences of clusters. Then, the system verifies some constraints for each crowd and labels them as unusual if necessary. Subsequently, one or more unusual crowds compose unusual events.

Let $DB_{CDR} = \{call_1, call_2, \dots, call_n\}$ denote the set of all calls collected from a mobile phone network. We define a call as a tuple $call_i = \langle t_i, v_j, l_k \rangle$, which means a user v_j makes or receives a call at location l_k at timestamp t_i , where $v_j \in V, t_i \in T, l_k \in L$. V is the set of all users and T denotes all possible timestamps. Specifically, l_k stands for the geographical location of the k^{th} mobile network antenna and L means the set of locations of all antennas found in DB_{CDR} . We define the individual mobility trajectory [10, 27] for each user as follows.

Definition 1. Individual Trajectory: A user v_j ’s mobility trajectory from start time t_p to end time t_q is defined as a sequence of spatio-temporal tuples $s_j^{pq} = \bigcup(t_i, l_k)$, where $t_p \leq t_i \leq t_q$ and $s_j^{pq} \in S$. S stands for the set of user trajectory sequences.

Cylindrical Cluster. The first step to identify crowd events from individual trajectories is to find, at any specific timestamp, clusters of individuals that are close in space. However, since CDR data is very sparse on the time scale (i.e., users do not make calls regularly), we propose the concept of *cylindrical cluster* in coarse-grained spatio-temporal data. Finer grain clustering methods, such as density-based clustering [8], cannot be applied as the antenna is the lowest level of spatial resolution available in the data. Indeed, users are already clustered by association to the antenna they use at each call (which defines a specific coverage area in the city, ranging from a few hundred squared meters to a few kilometers).

Definition 2. Cylindrical Cluster: Given a CDR database DB_{CDR} which contains individual calls with time and antenna information, and a scale threshold ϵ_n , the cylindrical cluster CC_t at timestamp t is a non-empty subset of users $V_t \subseteq V$ satisfying the following conditions:

- Connectivity. $\forall v_i \in V_t, v_i$ makes at least one call by using antenna a_x , in the interval $[t - \epsilon_t, t + \epsilon_t]$.
- Scale. The number of users $|V_t|$ in CC_t is no less than ϵ_n .

Figure 2(a) shows an illustrative example for cylindrical clusters. Given a timestamp t_1 , we can see that *user1, user2, user3* and *user4* make calls during time interval $[t_1 - \epsilon_t, t_1 + \epsilon_t]$. Also, *user3, user1* and *user2* use the same antenna which is different from *user4*’s. Then they are clustered into two groups. One potential issue is that there may exist multiple locations for one single user if she/he makes multiple calls during time interval $[t_1 - \epsilon_t, t_1 + \epsilon_t]$. A number of

methods can be considered to assign one single location from multiple locations, such as the central position or the most common position. We use the most common position due to its ease of calculation and understanding.

Crowd. In order to detect crowds lasting for a certain amount of time we need to consider shared characteristics between clusters detected in consecutive timestamps.

Definition 3. Crowd: Given a CDR database DB_{CDR} with individual trajectories, a lifetime threshold ϵ_{lt} , a consecutive intersection threshold ϵ_{ci} and a commitment probability threshold ϵ_p , a crowd C is a sequence of consecutive cylindrical clusters $\{CC_{t_m}, CC_{t_{m+1}}, \dots, CC_{t_n}\}$ which satisfy the following constraints:

- Movement. The number of total locations in one crowd is more than one.
- Durability. The lifetime of C , $C.lt$, namely the number of consecutive clusters, is greater than ϵ_{lt} , i.e., $C.lt \geq \epsilon_{lt}$ where $C.lt = n - m + 1$.
- Commitment. At least ϵ_{ci} users appear in each cylindrical cluster with existence probability ϵ_p .

The movement and durability characterizations specify the types of crowd we are interested in. The commitment instead characterizes the fact that a certain subset of users needs to participate to all clusters. Again, due to the spatio-temporal sparsity of the CDR data, the computation of the commitment of an user requires some further considerations. Therefore, we propose the concept of existence probability, which is designed to overcome CDR sparsity. Indeed, as an individual is not constantly making calls, consecutive timestamps could not see all users in the cluster making calls.

We design the existence probability of one user locating in a cluster at timestamp t as the proportion of the number of users in CC_t to the number of users in CC_{t-1} . The intuition for the definition of existence probability is that the user has conformity to follow others in the group that she or he was assigned to [24]. For example, the existence probability of *user3* in Figure 2(a) at time t_2 is $1/3$. In timestamp t_1 , *user1*, *user2*, and *user3* stay in cluster CC_{t_1} , and one of them, user *user2*, goes to cluster CC_{t_2} at timestamp t_2 . *user1* and *user3* do not make calls in timestamp t_2 , which results in the uncertainty of their locations. Thus, we assign them the probability to stay with *user2*, which is in cluster CC_{t_2} . Furthermore, we make the existence probability decay over time, i.e., if a user does not appear in consecutive timestamps, such as user *user4* in timestamp t_3 and t_4 . Her existence probabilities in Figure 2(a) are $[0, 1, \frac{1}{2}, \frac{1}{2} \times \frac{2}{3}]$ at each timestamp, respectively.

Considering that a crowd is a sequence of clusters, we use the standard terminology of sequential pattern mining and affirm that: a crowd C is called a closed crowd if it has no super crowds, which means there does not exist super sequences containing C .

Unusual Crowd. Usually, people have their own mobility trajectories in daily lives, such as going from home to work place everyday. When people go to attend a concert or a protest, their trajectories differ from their usual ones.




Fig. 2. Illustrative Examples of Cylindrical Cluster and Closed Crowd.

The definition of crowd given above includes both usual daily trajectories (e.g., commuting) as well as unusual event trajectories (e.g., protests). This is, for instance, what the method in [32] aims to do. As we will show in the experiments section, such method generates an enormous amount of events, as opposed to what a city would need in order to identify specific unusual events. Here we define the concept of mobility profile to capture people’s normal movement behaviors, by comparing with which we can detect abnormal mobility behaviors.

Definition 5. Mobility Profile: Given a CDR database DB_{CDR} with individual trajectories, one’s mobility profile is the groups of locations she/he visited for each time unit (hour) in every day. Notice that a location here corresponds to an antenna.

Definition 6. Unusual Crowd: Given the mobility profiles of users, a similarity threshold ϵ_{si} , a closed crowd C is said to be an unusual crowd UC if the average similarity between the trajectory of each user in the crowd and her/his mobility profile in corresponding time intervals is less than ϵ_{si} .

Unusual Event Detection. Due to the inaccuracy of CDR data and to the introduction of the existence probability concept, it is possible that two or more crowds share users and thus they represent the same event. Moreover, it is possible that many crowds might correspond to the same large event (e.g., two parades converging to the same square). To group together these unusual crowds, we define the concept of unusual event:

Definition 7. Unusual Event: Given two unusual crowds UC_i and UC_j , UC_i and UC_j are connected into one unusual event if they satisfy the following principles:

- Overlapping: The ending time $C_i.t_{end}$ of crowd C_i is temporally close to the beginning time $C_j.t_{begin}$ of other crowd C_j , w.r.t. $C_j.t_{begin} < C_i.t_{end}$.
- Sharing: The number of common users, $|C_i \cap C_j|$, is larger than or equal to half of the total users $|C_i \cup C_j|$.

An unusual event is a set of unusual crowds $E = \{UC_1, UC_2, \dots, UC_n\}$ in which any two unusual crowds are connected to each other by a path. Here one separate unusual crowd is also an unusual event, if it does not connect with others. Based on the discussed concepts above, we formalize the unusual event detection problem as follows.

Problem 1. Unusual Event Detection: Given all detected crowds during the interval of two timestamps, the goal of unusual event detection is to extract all unusual events happening in the time interval.

Unusual crowd event detection in mobile phone CDR data faces several unique challenges. First, the sparseness of CDR data comes from not only the fact that a user's location is recorded only when a call is made but also the way that this location is approximated as the cover area of an antenna that is being used by this call. To solve the temporal and spatial sparseness of CDR data, we propose to define user existence probability that can overcome the fact that a user's location is recorded only when a call is made, and also to leverage the idea of cylindrical cluster to address the coarseness of user locations as they are recorded as the cover area of involved antenna. Moreover, the problem is targeted at inferring unusual events rather than people daily routines. To achieve so, we propose the concept of mobility profile to distinguish unusual crowding behavior from daily movements.

3 Unusual Event Detection Framework

Given the formal definitions above, we describe now an innovative and efficient framework to detect unusual crowd events from CDR data. Our framework is composed of four parts: cylindrical cluster detection, closed crowd detection, unusual crowd detection, and unusual event detection.

Cylindrical Cluster Detection. Given the database of the individual calls with the respective time and antenna information, a duration threshold ϵ_t , and a scale threshold ϵ_n , the Cylindrical Cluster Detection algorithm maintains at each timestamp t the set of users observed from each antenna a , in the time interval $[t - \epsilon_t, t + \epsilon_t]$. Then, for each timestamp it returns all the set of users whose size is larger than ϵ_n . All the detected cylindrical clusters are stored in *ClusterDB*.

Closed Crowd Detection. The input for crowd detection is a set of cylindrical clusters *ClusterDB* extracted at each timestamp. There are three constraint thresholds considered in our crowd definition: *movement*, *durability*, *commitment*. Explicitly, if the subcrowd of one crowd meets the *durability* and *movement* constraints, it will satisfy the *commitment* constraint also. Thus the crowd definition satisfies the requirement of downward closure property, and then it is unnecessary to output all crowds, including the subcrowds of closed crowds. To avoid the redundancy resulted from outputting subcrowds, we can follow the Lemma 1 to decide if a crowd is closed or not.

Lemma 1. *A crowd C with clusters $\{CC_{t+m}, CC_{t+m+1}, \dots, CC_{t+n}\}$ is a closed crowd, if there does not exist CC_{t+m-1} or CC_{t+n+1} that can be added to crowd C such that a new crowd is formed.*

The restriction of closed crowd contains two conditions, one is that no suffixed cluster can be appended into it and the other is that no prefixed cluster can be

merged in its front. To discover closed crowds in cluster database at current timestamp t , the first condition is easy to check: if there exist clusters in next timestamp $t + 1$ that can be appended to current crowd C , then the process will continue; if not, we only need to verify whether current crowd C is the subcrowd of crowds formed at current timestamp t . It is not necessary to check every crowd at previous timestamps because that current crowd at timestamp t can only be the subcrowd of crowds ending at timestamp t .

Figure 2(b) shows an illustrative example for this process. Suppose that crowds C_1 and C_2 are found as closed crowds, if there is no cluster at timestamp t_6 that can be appended to crowd C_3 , then we need to further check whether it is the subcrowd of previous crowds. It is obvious that it is impossible for C_3 to be the subcrowd of crowds ending at t_4 or earlier timestamps, such as C_1 , but it is possible to be the subcrowd of crowds ended at t_5 , such as C_2 .

To find all closed crowds in *ClusterDB*, we start with iterating each timestamp in an increasing order. At each timestamp t , we check whether each candidate crowd at timestamp $t - 1$ can be appended by clusters at timestamp t . If the candidate crowd satisfies the *movement* and *durability* constraints, and at the same time it is not the subcrowd of crowds ending at timestamp $t - 1$, then we can output the current candidate as a closed crowd. The current candidate crowd can then be appended by one more cluster to form a new candidate crowd at t . The candidate crowd set contains all crowds which can be appended by a new cluster at t . Then we put all clusters at timestamp t to it to form a new candidate crowd set at t . This order of adding candidate crowd to candidate set guarantees that we only need to check whether the potential crowd is the subcrowd of closed crowds ending at the same timestamp.

Complexity: The extraction of closed crowds is similar to the extraction of closed frequent sequential patterns whose complexity in the worst case can be approximated with $\mathcal{O}(|A|^2 * |T|)$ where $|A|$ is the number of antennas (i.e. clusters) and $|T|$ is the number of timestamps.

Unusual Crowd Detection. With the detected closed crowds, we further verify whether their users present unusual or regular behaviors. As introduced in Section 2, we use mobility profile to decide whether users' movement trajectories are unusual.

To generate the mobility profiles, we scan the historic CDR data once to record the specific locations a user visited at each timestamp during every time period. For example, *user4*'s existence probability vector in corresponding crowd is $\mathbf{w}_c = [0, 1, \frac{1}{2}, \frac{1}{2} \times \frac{2}{3}]$ in Figure 2(a). His profile vector is extracted from his mobility profile at corresponding timestamps (from t_1 to t_4), i.e. $\mathbf{w}_m = [\frac{0}{3}, \frac{2}{2+5}, \frac{1}{4+1}, \frac{2}{2+1}]$. There are several ways to define the similarity between user's mobility profile and his trajectory in the crowd. We use cosine similarity to calculate the similarity score, because of its ease of understanding and implementation. The cosine similarity between two vectors \mathbf{w}_c and \mathbf{w}_m is defined as:
$$CosSim(\mathbf{w}_c, \mathbf{w}_m) = \frac{\mathbf{w}_c \cdot \mathbf{w}_m}{\|\mathbf{w}_c\| \|\mathbf{w}_m\|}.$$

The Unusual Crowd Detection algorithm first calculates for each user in the crowd the similarity between her trajectory and her own mobility profile.

Then the similarities obtained are averaged, and the obtained value is greater than the *similarity* parameter ϵ_{si} , it is an unusual crowd.

Complexity: The mobility profile construction requires a scan of the dataset, therefore its complexity is $\mathcal{O}(DB_{CDR})$. The detection of Unusual Crowds requires for each crowd the computation of the cosine similarity for all the users being part of a crowd, thus its complexity is $\mathcal{O}(|C|*|V|)$ where $|C|$ is the number of crowds and $|V|$ the number of users.

Unusual Event Detection. With discovered unusual crowds, we finally detect their relationships and connect them into one event if they meet the requirements of Definition 7. In this step, we use graph theory to find and generate unusual events. First if two unusual crowds satisfy both *overlapping* and *sharing* principles, we create an edge to connect them. With this generated graph, where each node is one unusual crowd and an edge indicates that two crowds belong to the same event, the event detection is to generate all components in the graph. Note that this graph may not only be disjoint but also include single nodes. Each component or single node is an unusual event that is our final goal of this work. The first part of this algorithm checks if two unusual crowds can be connected to each other by parameters *overlapping* and *sharing*. The second part generates all the components in the unusual crowds graph, where any graph algorithm can be used. The detected event contains the users in each cluster and its corresponding timestamp and location.

Complexity: The detection of Unusual Events requires a pair-wise comparison between all the Unusual Crowds, therefore the complexity of this procedure is $\mathcal{O}(|UC|^2)$ where $|UC|$ is the number of Unusual Crowds.

4 Experiments

4.1 Experimental Setup

CDR Data. The D4D Orange challenge made available data collected in Cote d'Ivoire over a five-month period, from December 2011 to April 2012. The datasets describe call activities of 50,000 users chosen randomly from every two weeks. Specifically, the data contains the cell phone tower and a timestamp at which the user sent or received a text message or a call in the form of tuple $\langle \text{UserID}, \text{Day}, \text{Time}, \text{Antenna} \rangle$. Each antenna is associated with location information. To avoid privacy issues, the data has been anonymized by D4D data provider.

From the CDR data, we find that about 63% users do not make calls in consecutive hours and 19% users make calls in only two consecutive hours. The pattern demonstrates the necessity of existence probability for user's location estimation, as most of users do not make regular and consecutive calls at each timestamp. We also observe that the probability that there is one hour between one user's two calls is more than 75% and that is 8% for a two-hour interval. In total, there are more than 80% two consecutive calls whose intervals are at most two hours. These observations demonstrate the challenges of CDR data's

Table 1. Comparison of the unusual event detection (UE) and gathering detection (GAT) [32].

Period	Date	Event Name [22]	UE	#UE	GAT	#GAT
Dec. 05 - Dec. 18	Dec. 07	Anniversary of Felix Death	✓	20	✓	287
	Dec. 11	Parliament election	✓		×	
	Dec. 17	Violence	✓		✓	
Dec. 19 - Jan. 01	Dec. 25	Christmas day	×	36	×	56
	Dec. 31	New year eve	✓		✓	
	Jan. 1	New year day	✓		×	
Jan. 02 - Jan. 16	Jan. 08	Baptism of Lord Jesus	✓	31	✓	176
	Jan. 14	Arbeen Iman Hussain	✓		×	
Jan. 17 - Jan. 29	Jan. 17	Visit of Hilary Clinton	✓	15	✓	481
	Jan. 18	Visit of Kofi Annan	✓		✓	
Jan. 30 - Feb. 12	Jan. 30	ACNF 2012 vs Angola	✓	58	✓	310
	Feb. 04	ACNF 2012 vs Equatorial Guinea	✓		✓	
	Feb. 04	Mawlid an Nabi Sunni	✓		✓	
	Feb. 05	Yam	✓		×	
	Feb. 08	ACNF 2012 Semi Final VS. Mali	×		✓	
	Feb. 09	Mawlid an Nabi Shia	✓		✓	
Feb. 13 - Feb. 26	Feb. 13	Post African Cup of Nations Recovery	✓	52	✓	152
	Feb. 22	Ash Wednesday	✓		✓	
Feb. 27 - Mar. 10	None			26		269
Mar. 11 - Mar 25	Mar. 12	Election of National Assembly President	✓	17	✓	342
	Mar. 13	Election of National Prime Minister	✓		✓	
Mar. 26 - Apr 08	Apr. 01-04	Education International Congress	✓	75	✓	1220
	Apr. 06	Good Friday	✓		✓	
Apr. 09 - Apr. 22	Apr. 09	Easter Monday	✓	10	✓	33
	Apr. 13-14	Assine fashion days	✓		✓	
Total			23/25	340	19/25	3326
Precision			0.0676		0.0057	
Recall			0.9200		0.7600	

spatio-temporal sparseness, which makes the design for degenerative existence probability reasonable for the coarse-grained CDR data.

Comparison Methods. To the best of our knowledge, this is the first work to detect unusual crowds and events in spatio-temporal data, and it is also the first time that we discover moving clusters in CDR data. We compare the results of our approach with GAT described in [32] and MOV in [17], as the methods employed in these work are also able to identify moving crowds. However, those methods are not designed to work on CDR and have to be adapted to perform the comparison. GAT defines a method to detect the gatherings in a trajectory dataset. A gathering is a sequence of spatial clusters with a certain number of committed users being member of an enough number of clusters. We use the same setting with GAT for parameters that indicate the same physical meanings in both methods. Clearly by following our intuition and goal of problem design, there should not exist any crowd or event at most days. Based on the results of parameter analysis in Section 4.3 and the developed Visual Analytics System, we selected the following parameters $\epsilon_n=20$, $\epsilon_{lt}=4$, $\epsilon_{ci}=10$, $\epsilon_p=0.2$, and $\epsilon_{si}=0.2$. Since they correspond to a probability to find unusual crowds to be around 10-15%, which helps us focus on rare events (as opposed to business as usual events).

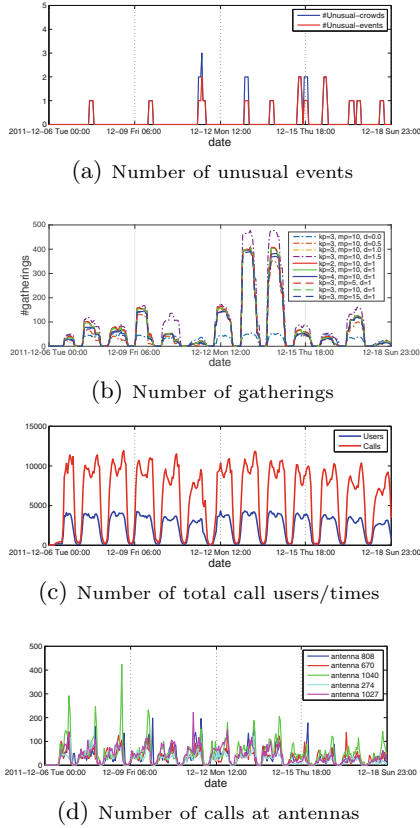


Fig. 3. Time series of unusual events, gatherings, users, and calls in the first two-week.

4.2 Experimental Results

Detected Unusual events. Table 1 reports a series of events occurred in Abidjan in the different periods covered by the datasets. In order to perform a fair study of the effectiveness of our method in comparison with GAT, we selected a third part set of events reported in [22]. To limit the explosion in the number of detected gatherings, we have set the most restrictive values for the remaining parameters: $d = 0.0$, $k_p = 2$, and $m_p = 5$. Moreover, we report the total number of generated events by both methods for each two-week period and subsequently generate *Precision* and *Recall* scores for both algorithms. It is possible to notice that our method detects a lower number of unusual events w.r.t. GAT. This is reflected in a higher value of Precision. Although, our method reports a lower number of events, it is able to detect a greater number of ground truth events, and this corresponds to a higher value of Recall. Notice that the two measures represent an estimation of precision and recall since the ground truth is not given. Indeed the list of events in [22] is not comprehensive of all events that happened in Ivory Coast in the monitored 5-month period, and this explains the low precision of both methods. This is the reason why we did not try

to find the optimal values of the parameters to maximize Precision and Recall, but instead set such values based on the general criteria to find unusual crowds only in around 10-15% of the hours. However, the list in Table 1 gives a good basis for comparison and shows that our method is 10 times more precise than GAT.

We further perform comparisons with MOV, where the authors introduce the concept of moving clusters. Extracting moving clusters is equivalent to run our method with the parameter ϵ_p (the probability of a user to be committed) and ϵ_{si} (the similarity threshold between the mobility profile and the trajectories in the event) to 1. With these parameter settings the algorithm was not able to find any moving clusters. This is due to the fact that our method is able to handle the spatial and temporal sparsity of CDR data, while the MOV method is designed to work with GPS trajectories.

Time-Series. We report the time series of the numbers of unusual events and gatherings detected with different input parameter settings by using our method and the GAT algorithm in Figures 3(a) and 3(b). We try to match the same parameters we used in our method. For the minimum lifetime as well as the minimum number of objects that should belong to a cluster, we choose the same values adopted in the study of the effectiveness of our method ($\epsilon_{lt} = 4$ and $\epsilon_n = 20$). The rest of the input parameters of the algorithm to detect gatherings are the following: d is the minimum distance necessary to connect clusters detected in two consecutive time snapshots; k_p is the minimum number of time snapshots required to consider an user as a participant; m_p is the minimum number of participants to create a gathering. For these input parameters, we tried different enumerations to span the full admissible ranges. As it is possible to see, the number of detected gatherings is very high even if the parameters are chosen to be very restrictive. All the graphs show a daily trend, demonstrating that this method is not able to find unusual events our proposed method. Indeed, GAT can detect a large number of gatherings every day, which might not correspond to specific unusual events.

To further evaluate our discoveries, we check the total communication volumes and the specific antenna activities. We can clearly see that between Dec. 06 and Dec. 18, in Figure 3(c) there exist periodic patterns on each day without obvious peak values corresponding to the discovery of crowds—**anniversary of Felix death** on Dec. 07 and **Parliament election** on Dec. 11. Furthermore, the events on the day of parliament election involved five antennas. Their communication activities are plotted in Figure 3(d). Obviously there do not exist correlations between corresponding antenna activities and our unusual crowd/event output. These two regular and stable time series of communication activities further confirm the effectiveness of our problem design. These examples show that detecting unusual events is a complex task, which cannot be easily accomplished by looking at outliers in call time series. Thus, methods like [6, 21] are not directly applicable.

Spatial Distributions. Another comparison performed against GAT regards the spatial distribution of the detected events/gatherings. For both methods, we




Fig. 4. The unusual events (a) by our methods and the gatherings (b) by GAT detected on December 11th. Colors range from green to red as function of the number of detected participants

select the results obtained on December 11th. For the GAT method we select the results with lowest number of detected gatherings. In Figures 4(a) and 4(b) we report the detected events and detected gatherings respectively. Notice that a Voronoi tessellation has been applied in order to associate a covering area to each antenna. Our method detect 2 events, *Event 1* (left) covers 3 antennas and lasts for 4 hours. *Event 2* (right) covers 2 antennas and also lasts for 4 hours. For the same day, the GAT algorithm detects many gatherings (25) occurring in different regions of the city. This is probably due to the fact that the typical mobility profiles of the users are not taken into account in the process and thus recurring and unusual events are both detected. Moreover, if we consider the lifetime of gatherings occurring in the same locations of our events, we notice that it is generally longer. For example, a gathering, covering the same antennas of *Event 1*, lasts for 14 hours. Another characteristic of the gatherings is that they happen in the same location at different times. Instead, in our model, we define a method to consider those as one large event. In summary, with our method it is possible to identify events that occur occasionally in a precise zone of the city and happen in a precise period of time, while the other method detects several events without any distinctions between the periodical and the unusual ones.

4.3 Efficiency and Parameters

Our algorithms are implemented in Python 2.7.5, and all experiments were performed on a laptop running Windows 7 with Intel(R) Core(TM) i7-2720QM CPU@2.20GHz (2 cores) and 8GB memory. All related experiments are running on the first two-week dataset, which contains about two million CDR historic data. We simulate each experiment with specific parameter setting for 100 times to get both the average running time and standard deviation. In general, the algorithms for detecting unusual crowd events are efficient, in the fact that it only takes about 30 seconds to two minutes on two million CDR data. Furthermore, the execution of our methods is stable among different runs.

We evaluate the effect of parameter setting on the number of detected unusual crowds and discuss the guidelines for determining parameter settings. We find that the algorithm is particularly sensitive to ϵ_{lt} , ϵ_p , and ϵ_{si} . ϵ_{lt} is indicative of



Fig. 5. Real-time view of city map and statistics.

the duration of moving crowds. Based on the definition of *commitment*, a larger ϵ_p can produce more compact crowds, which have much higher probabilities to be unusual events. Finally, the lower *similarity* ϵ_{si} threshold between regular mobility profiles and specific trajectory we set, the more crowds will consist of people whose mobility behavior differs from their typical profiles.

We would like to point out that there is not an unique way to optimally select the values of the parameters, as this strongly depends on the end-user application. For instance, if a city manager is interested in monitoring visitors to a museums, she might set different values of ϵ_{lt} and ϵ_p , compared to the monitoring on a protest. We develop a visual analytics tool described in Section 5 to help end users explore and test the detected results under different parameter settings for different applications.

5 Visual Analytics Prototype System

We have developed a visual analytics system to support the exploration of the unusual crowd events based on the proposed framework. The system allows end users—such as analysts and city managers—to analyze the formation and evolution of crowds, and study the impact of different parameters on the obtained results, heuristically suggesting possible changes to get more meaningful results depending on the desired application. The interface consists mainly of two components: the map overview of the observed city (Figure 5 (a)) and the statistics of users, crowds, and events (Figure 5 (b)).

Map View. In the map, the system visualizes the latest clusters, crowds and unusual events detected in the form of polygons as shown in Figure 5 (a).



Fig. 6. Analyst statistics view

The polygons are the convex hulls of the location updates of the users belonging to one of the aforementioned groups. The clusters detected by the system at a given timestamp are visualized on the map as green polygons. On mouse-over, the UI shows a pop-up window with the cluster attributes, including 1) timestamp—the timestamp when the cluster was detected, 2) #users—the number of users that are a part of the cluster, 3) area—the area covered by the Cluster polygon in square kilometers, 4) density—the ratio between the number of users in the Cluster and its area, and 5) POIs—the list of Points Of Interest located within the cluster area. The detected Crowds and the Unusual events are visualized on the map as blue and red polygons, respectively. Similar to the cluster visualization, a set of properties can be shown in a pop-up window.

Statistics View. We envision the tool to be used by two different types of actors: the Analyst, which is in charge of setting up the analytics system, and the city manager, who has to take actions based on the identified unusual crowd events. The City manager tab contains a single time-series graph representing the number of Clusters, Crowds, Unusual Events detected at every timestamp as displayed in Figure 5 (b). This tab contains the most crucial outcome of the analytics performed and it provides an intuitive way to represent the most recent mobility patterns of the city. The Analyst tab contains a richer set of statistics in Figure 6. In order to make efficient use of available space we fit the graphs into collapsible panels into groups of semantically relevant statistics, including 1) Cumulative—the cumulative trends of the detected clusters, crowds and unusual events, 2) Detection per minute—the time-series of the number of clusters, candidate crowds, crowds, and unusual events, 3) Event monitoring—the time-series of maximum and minimum value of lifetime, number of committed users, total number of users, and similarity of the candidate crowds, 4) Cluster monitoring—the time-series of the maximum size of the detected clusters, and their minimum spatial radius. In addition, a red dashed-line corresponding to each parameter is shown to depict parameter efficacy (e.g. event monitoring,

cluster monitoring). This allows the Analyst to understand the role of each parameter on the obtained results and to set the most appropriate parameters for the specific application in scope.

6 Related Work

The availability of mobility data has offered researchers the opportunity to analyze both individual's and group's moving behaviors. In [12,14], the authors defined a methodology to extract dense areas in spatio-temporal databases, thus identifying where and when dense areas of mobile objects appear. A similar definition was proposed in [17], where the authors introduced the concept of *moving clusters*. Following these ideas, several group and cluster mobility pattern models have been proposed [1,16,19,31]. For instance, in [2,11,28,29] the concept of *flock* is widely investigated. A flock is a group of objects that travel together within a disc of some user-specified size for at least k consecutive timestamps. The main limit of this model is that a simple circular shape does not reflect natural grouping in reality. A marginal different concept *convoy* was introduced in [15,16], where a density-based clustering is adopted instead of the radius of the disk. Li et al. [18] proposed a more general type of trajectory patterns: *swarm*. The swarm is a cluster of moving objects that lasts for at least k timestamps, possibly non-consecutive. Another group mobility model was introduced in [32], called *gathering*, whose novelty regards the introduction of the concept of commitment.

Other interesting works dealing with the detection of anomalies in city traffic flow are presented in [6,21,25]. In [21], the authors use likelihood ratio test statistic (LRT) on GPS trajectories of taxis to detect traffic flow anomalies. Transportation model detection problem is studied in mobile phone data and GIS data [23]. A passive route sensing framework is introduced to monitor users' significant driving routes with low-power sensors in mobile phones [20]. However, these works do not address the problem detecting unusual events considering people mobility but are more focused on traffic flow analysis through an aggregation of the information. On the contrary, in this paper we are interested in detecting such events that involve a large number of people whose current mobility differs from their typical one.

All the above works are designed and tested on high-resolution trajectory data, such as the one provided by GPS systems. Low-resolution location data collected from telecommunication operators, on the other hand, is much more pervasive resulting in a much larger sample of the population being monitored, see [4,5,9,13]. In this paper, we propose a new method to mine coarse grain mobile phone data (in the form of CDR) to detect unusual crowd events. Indeed, the aim of our work is to detect events that involve a large number of people performing unusual activities. To do so, we compute a similarity between the mobility profile of the users and their trajectories in group pattern. This extends, thus, the concept of commitment since users need to be committed and have trajectories that differ from the ones in their mobility profiles. Our method

however is able to identify moving events that span several locations over time, and involve a subset of committed users, something that could not be detected by using the methods in [26,30].

7 Conclusion and Future Work

In this paper, we formally define the problem of inferring unusual crowd events from mobility data. Previous work on event detection is limited on inferring the usual event from the fine-grained GPS data. Our problem definition differs by characterizing the unusual crowd events and we present a new methodology to extract them from coarse-grained CDR data. The main contributions of this paper w.r.t. existing methods are the ability to analyze temporally and spatially sparse data as CDRs and the definition of a subclass of events which are unusual to its attendees. Our experimental results demonstrate the effectiveness of our proposed method in real-world mobile datasets.

Despite the promising results of the present work, there is still much room left for future work. First, while this proposed method relies on Visual Analytics to help end users set parameters, we are planning to design algorithms to determine parameters for specific applications of interest as well as an optimization procedure for evaluation metrics. Moreover, we are working toward combining mobile and social media data together to detect unusual events. In doing so, we have the potential to detect and monitor crowding activities in real time, and eventually yield a better and smarter planet.

References

1. Appice, A., Malerba, D.: Leveraging the power of local spatial autocorrelation in geophysical interpolative clustering. *DMKD* **28**(5–6), 1266–1313 (2014)
2. Benkert, M., Gudmundsson, J., Hbner, F., Wolle, T.: Reporting flock patterns. *Computational Geometry* **41**(3), 111–125 (2008)
3. Calabrese, F., Di Lorenzo, G., McArdle, G., Pinelli, F., Van Lierde, E.: Real-time social event analytics. In: *Netmob 2015* (2015)
4. Calabrese, F., Ferrari, L., Blondel, V.: Urban sensing using mobile phones network data: a survey of research. *ACM Comput. Surv.* (2014)
5. Calabrese, F., Pereira, F.C., Di Lorenzo, G., Liu, L., Ratti, C.: The geography of taste: analyzing cell-phone mobility and social events. In: Floréen, P., Krüger, A., Spasojevic, M. (eds.) *Pervasive 2010*. LNCS, vol. 6030, pp. 22–37. Springer, Heidelberg (2010)
6. Chawla, S., Zheng, Y., Hu, J.: Inferring the root cause in road traffic anomalies. In: *IEEE ICDM 2012*, pp. 141–150 (2012)
7. Dong, Y., Yang, Y., Tang, J., Yang, Y., Chawla, N.V.: Inferring user demographics and social strategies in mobile social networks. In: *KDD 2014*, pp. 15–24. ACM (2014)
8. Ester, M., Peter Kriegel, H., Sander, J., Xu, X.: A density-based algorithm for discovering clusters in large spatial databases with noise. In: *ACM SIGKDD 1996*, pp. 226–231 (1996)

9. Georgiev, P., Noulas, A., Mascolo, C.: The call of the crowd: Event participation in location-based social services. In ICWSM 2014 (2014)
10. Giannotti, F., Nanni, M., Pinelli, F., Pedreschi, D.: Trajectory pattern mining. In: ACM SIGKDD 2007, pp. 330–339. ACM, New York (2007)
11. Gudmundsson, J., van Kreveld, M.: Computing longest duration flocks in trajectory data. In: ACM GIS 2006, pp. 35–42. ACM, New York (2006)
12. Hadjieleftheriou, M., Kollios, G., Gunopulos, D., Tsotras, V.J.: On-line discovery of dense areas in spatio-temporal databases. In: Proc. of the 7th International Conference on Advances in Spatial and Temporal Databases, SSTD 2003, pp. 306–324 (2003)
13. Isaacman, S., Becker, R., Cáceres, R., Martonosi, M., Rowland, J., Varshavsky, A., Willinger, W.: Human mobility modeling at metropolitan scales. In: Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services, MobiSys 2012, pp. 239–252. ACM, New York (2012)
14. Jensen, C.S., Lin, D., Ooi, B.C., Zhang, R.: Effective density queries on continuously moving objects. In: IEEE ICDE 2006, pp. 71–81. Washington, DC, USA (2006)
15. Jeung, H., Shen, H.T., Zhou, X.: Convoy queries in spatio-temporal databases. In: IEEE ICDE 2008, pp. 1457–1459. IEEE Computer Society, Washington (2008)
16. Jeung, H., Yiu, M.L., Zhou, X., Jensen, C.S., Shen, H.T.: Discovery of convoys in trajectory databases. Proc. VLDB Endow. **1**(1), 1068–1080 (2008)
17. Kalnis, P., Mamoulis, N., Bakiras, S.: On discovering moving clusters in spatio-temporal data. In: Medeiros, C.B., Egenhofer, M., Bertino, E. (eds.) SSTD 2005. LNCS, vol. 3633, pp. 364–381. Springer, Heidelberg (2005)
18. Li, Z., Ding, B., Han, J., Kays, R.: Swarm: mining relaxed temporal moving object clusters. Proc. VLDB Endow. **3**(1–2), 723–734 (2010)
19. Lu, E.H.-C., Tseng, V.S., Yu, P.S.: Mining cluster-based temporal mobile sequential patterns in location-based service environments. IEEE Trans. on Knowl. and Data Eng. **23**(6), 914–927 (2011)
20. Nawaz, S., Mascolo, C.: Mining users’ significant driving routes with low-power sensors. In: ACM SenSys 2014, pp. 236–250. ACM (2014)
21. Pang, L.X., Chawla, S., Liu, W., Zheng, Y.: On detection of emerging anomalous traffic patterns using gps data. Data Knowl. Eng. **87**, 357–373 (2013)
22. Paraskevopoulos, P., Dinh, T.-C., Dashdorj, Z., Palpanas, T., Serafini, L.: Identification and characterization of human behavior patterns from mobile phone data. In: International Conference on the Analysis of Mobile Phone Datasets (NetMob 2013) (2013)
23. Stenneth, L., Wolfson, O., Yu, P.S., Xu, B.: Transportation mode detection using mobile phones and gis information. In: GIS 2011, pp. 54–63. ACM, New York (2011)
24. Tang, J., Wu, S., Sun, J.: Confluence: conformity influence in large social networks. In: ACM SIGKDD 2013, pp. 347–355. ACM, New York (2013)
25. Telang, A., Deepak, P., Joshi, S., Deshpande, P., Rajendran, R.: Detecting localized homogeneous anomalies over spatio-temporal data. DMKD **28**(5–6), 1480–1502 (2014)
26. Traag, V.A., Browet, A., Calabrese, F., Morlot, F.: Social event detection in massive mobile phone data using probabilistic location inference. In: SocialCom 2011, pp. 625–628. IEEE (2011)
27. Trasarti, R., Pinelli, F., Nanni, M., Giannotti, F.: Mining mobility user profiles for car pooling. In: ACM SIGKDD 2011, pp. 1190–1198. ACM, New York (2011)

28. Vieira, M.R., Bakalov, P., Tsotras, V.J.: On-line discovery of flock patterns in spatio-temporal data. In: ACM GIS 2009, pp. 286–295. ACM, New York (2009)
29. Wachowicz, M., Ong, R., Renso, C., Nanni, M.: Finding moving flock patterns among pedestrians through collective coherence. *International Journal of Geographical Information Science* **25**(11), 1849–1864 (2011)
30. Witayangkurn, A., Horanont, T., Sekimoto, Y., Shibasaki, R.: Anomalous event detection on large-scale gps data from mobile phones using hidden markov model and cloud platform. In: ACM UbiComp 2013 Adjunct, pp. 1219–1228. ACM, New York (2013)
31. Wu, M., Jermaine, C., Ranka, S., Song, X., Gums, J.: A model-agnostic framework for fast spatial anomaly detection. *ACM Trans. Knowl. Discov. Data* **4**(4), 1–30 (2010)
32. Zheng, K., Zheng, Y., Yuan, N.J., Shang, S.: On discovery of gathering patterns from trajectories. In: IEEE ICDE 2013, pp. 242–253, Washington, DC, USA (2013)