Technology and Health Care 18 (2010) 275–284
DOI 10.3233/THC-2010-0590
IOS Press

Towards personalised ambient monitoring of mental health via mobile technologies

Pawel A. Prociow* and John A. Crowe
Electrical Systems and Optics Research Division, Faculty of Engineering, University of Nottingham, Nottingham, UK

Received 25 March 2010
Revised /Accepted 4 May 2010

Abstract. Managing bipolar disorder is an important health issue that can strongly affect the patient's quality of life during occurrences of depressive or manic episodes and is therefore a growing burden to healthcare systems. A widely practised method of monitoring the course of the disorder is by mood and general mental health questionnaires, which are nowadays often implemented on mobile electronic devices.
Detecting changes to daily routine and behaviour is of crucial importance as they can be symptomatic of an ongoing episode, or in the case of an external cause, may trigger such an episode.
Current mobile phones and geospatial technology provide a means of monitoring aspects of daily routine and lifestyle which may be valuable in facilitating self-management of the condition.
This manuscript introduces a methodology for analysing data obtained from a simple wearable system based on a mid-range mobile phone, along with trial results from a control group of three participants with no history of Bipolar Disorder. It is suggested that such an approach offers an unobtrusive, acceptable and low cost means of monitoring bipolar disorder patients that could significantly improve their care.

Keywords: Bipolar disorder, significant location discovery, pervasive monitoring, mobile psychiatry, Bluetooth encounters

1. Background

Bipolar disorder (BD) is a condition that affects almost 2% of the European population [16]. Sufferers of this disorder typically will go through several episodes of manic or depressive behaviour during their lifetime, often resulting in changes to their lifestyle that can occur within weeks or days [6,12]. It is the potentially serious consequences of such behavioural changes that can make the condition so debilitating and require long periods of carefully managed clinical care. A well-developed self-awareness is an important factor in the management of this condition with early recognition of symptoms being particularly important. Sufferers may use diaries and Personal Digital Assistants (PDA) as an aid in this process [2,12].

The types of lifestyle change, which may occur within a few weeks or even days, depending on the disorder type, that are associated with bipolar disorder, could include an increase or decrease in social interaction, both at work and during the individual's personal time [6,10]. The premise is that these

*Corresponding author. E-mail: prociow@ieee.org.

changes in mood and behaviour will be reflected in daily, and other regular routines and that this will be noticeable in both the locations that are visited and the individuals that are interacted with. What is more, evidence suggests that disruptions to daily routine, including social interactions, caused by external events may trigger an onset of an episode in bipolar patients [5,10]. Therefore, a successful 'early warning' system should be able to detect changes to those routines with monitoring of visited locations and social encounters of particular importance. The former can be monitored using the Global Positioning System (GPS), the latter via recording nearby Bluetooth devices [3,9].

The concept of personalised ambient monitoring (PAM) is to use technology to enhance a patient's self-awareness. In principle such a system would help patients to monitor their 'activity', which is linked to their mental state [6,10], whilst also providing a sufferer's support network with accurate and informative data at a mutually agreed level (since there are clearly privacy issues with such an arrangement).

The key elements for PAM are: a flexible sensor network providing data about patients' activities and behaviour; algorithms to process such data and detect changes in an individual's behavioural patterns and software to integrate these elements into a flexible platform with the capability of alerting not only the patient themselves of unusual behaviours but also their carers and clinicians. In addition there is a need to understand how the detected changes map onto an individual patient's mental status in order for the overall system to have the necessary specificity and sensitivity for it to be of clinical value [8].

An appropriate 'hub' for such a monitoring system is a mobile phone, since these are commonly carried everywhere and provide a communications link along with data storage and processing capabilities. Many mobile phones also provide location services using a built-in GPS receiver and most come equipped with a Bluetooth communications link enabling them to interact with other devices. These attributes provide a mechanism to acquire, log and process geospatial data collected by the user's phone.

This paper describes a preliminary investigation into the possibility of using the location aware functionality and Bluetooth radio of a mobile phone as a means of assessing behavioural changes that could be a result of a mood disorder.

2. Methods

2.1. GPS location

The Global Positioning System (GPS) is currently the most convenient method of obtaining geospatial data. Cheap commercial receivers able to provide location data to within an accuracy of around 3m [7] are widely available, and increasingly incorporated into many phones. Small 'key ring' devices with Bluetooth capability provide any Bluetooth enabled phone with access to positional information and this is the method employed in this study.

Of course, accurate determination of position is dependent upon the number of GPS satellites 'visible' to the receiver and the quality of data is often compromised for example in buildings and dense urban environments. The problem is mitigated to some degree, since data is provided on the number of visible satellites and if this is fewer than four the data is known to be unreliable [7].

2.2. Bluetooth encounters

In addition to collecting data from the GPS receiver the application was also programmed to scan for Bluetooth devices every 10 minutes. This provided a list of the unique addresses of nearby Bluetooth devices that were set to be discoverable. The proposed purpose of the scans was to provide both location information, via place-related infrastructure such as Bluetooth enabled office computers, as well as social interactions via detection of individuals' Bluetooth enabled phones within the scan boundary [3,9].

2.3. Experimental setup

A custom application was developed for the purpose of the trial. In order to maximise compatibility, the suite was created using the Java mobile device interface as the majority of modern phones are Java-enabled. The program could obtain position information either from a GPS receiver via a Bluetooth connection or from the internal GPS device (if present) utilizing the Java location services software interface. Regular scans for Bluetooth enabled devices were also performed in the main program flow.

The application was installed on a mid-range phone (Sony Ericsson G502 or Nokia 6120) and paired with a Bluetooth enabled GPS receiver (BlueNext BN-906GR). The premise of such an arrangement, over the use of an in-built GPS receiver, was the higher reliability of a quality external receiver in comparison with an internal one. What is more, the combined average battery life of a phone and an external receiver is higher than that of a phone with an internal GPS receiver.

The position data was recorded as sets of four values: latitude, longitude, speed and satellite count with this data sampled and stored every four seconds on the phone's memory card.

2.4. Trial description

The aim of the preliminary trial presented here was to validate the experimental setup and its potential to monitor the desired aspects of daily routine. The investigation was conducted on a small control group to enable data acquisition and analysis methods to be established that would be necessary for a trial involving patient groups.

The control experiment consisted of three male subjects who have no history of bipolar disorder and lasted for a period of time between three and four weeks in different locations within the UK. During the trial, subjects were asked to keep both the trial phone and the GPS receiver with them at all times, although they could either disable or terminate the monitoring application at any time.

Ethical permission to perform the trial was obtained from the University of Nottingham, Faculty of Engineering's Ethical Review Committee. The purpose of the study was described to the participants and informed consent to participate was taken, prior to commencing.

The data presented in this article was downloaded from the phone and stored on a secure server with access available only to the investigators. Only one of the subjects had the Bluetooth encounter capability included in the monitoring application.

2.5. Clustering

GPS based location data was expected to exhibit clustering around places where the subjects had spent a significant amount of time and therefore a suitable algorithm was required to analyse the data. There are several existing algorithms that facilitate the identification of clusters with the Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise algorithm (DBSCAN) having demonstrated its effectiveness in clustering geospatial records [1,13,17]. DBSCAN was developed at the University of Munich in 1996 and has since become one of the most commonly used algorithms for clustering large data sets.

The algorithm utilises the concept of points being density reachable from one another, which occurs when the entire distance between the points in question is continuously and densely populated by other points with the set. The DBSCAN clustering method is based on this measure, as opposed to the more conventional way of examining the neighbourhood of each point. The method requires only one

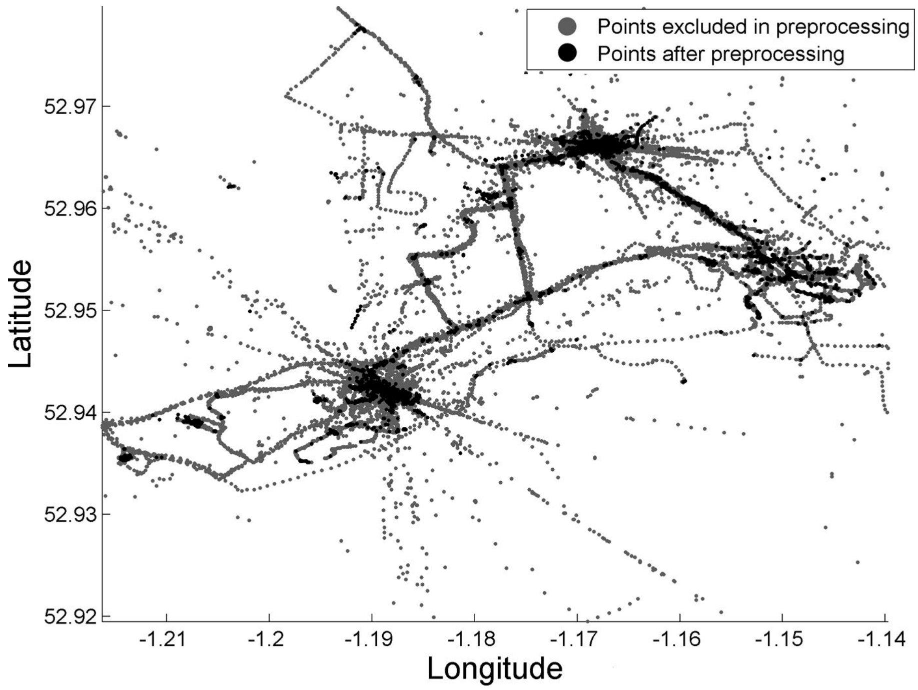

Fig. 1. Example dataset showing the effect of pre-processing the location data.

parameter which is the minimum number of points in the point's Epsilon-neighbourhood i.e. within a certain small distance [4].

There are several reasons favouring this method. Firstly, a small set of parameters makes it easy to generalize between subjects. Secondly, the method does not aim to cluster all the points in the input database meaning that a substantial subset may emerge un-clustered and so can be considered as "noise". This is a crucial feature as many points recorded will clearly not belong to a significant cluster (i.e. a meaningful location). For example there may be readings taken during the transition between locations and so are better left un-clustered.

2.6. Pre-clustering data processing

Pre-clustering processing of the GPS data consisted of removing all readings based on three or fewer satellites, which includes those producing the linear features in the bottom left quadrant of the Fig. 1. Since the main aim of the clustering process was to detect significant locations rather than map the journeys between them another pre-clustering step was to eliminate points where the recorded speed (provided by the GPS receiver using its internal calculations) suggested the subject was moving. Taking this into consideration a dataset was produced consisting of points gathered when the speed was less then 1km/h. Figure 1 shows data from one of the subjects distinguishing between points excluded in pre-processing and points used in further clustering basing on their reliability (i.e. determined using more than three satellites) and recorded speed.

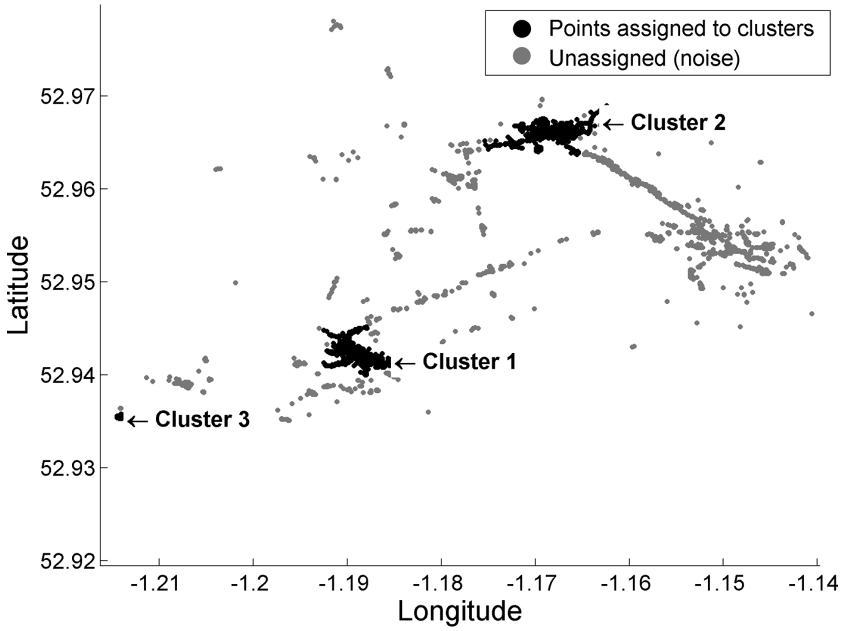

Fig. 2. Clustering with DBSCAN. Dataset from Fig. 1 after clustering with the MinPts parameter set to 5000.

3. Results

The three independent pre-processed databases contain results from three to four weeks of monitoring. Within this article, results from the single dataset that also contains Bluetooth encounter data are presented, as these are representative of the findings from the other two sets.

3.1. Clustering

After applying the pre-clustering steps, the dataset was subjected to clustering using the DBSCAN method. As mentioned, the algorithm takes one parameter, the minimum number of points per Epsilon-neighbourhood of a point (i.e. density). This parameter will be referred to as the minimum points (MinPts) parameter. The Epsilon value was set at 0.0005 deg in coordinate space which corresponds to an area of approximately 30 by 50 meters in the British Isles.

The MinPts parameter can be related to the total time spent by the subject in a particular location. In the presented approach data was first clustered with a parameter value of 5000. Taking into account the sampling period (one point is taken every 4s.) this corresponds to approximately 5.5 hours worth of records within a cluster throughout the entire dataset, although the actual time may be greater due to the fact that points with poor satellite coverage (mostly from indoor monitoring) were disregarded. Figure 2 shows results of such clustering performed on data presented in the previous figure.

The process of clustering a location can, depending on point density, be also classified as noise. This indicates that the spatial density is not high enough to constitute a cluster with the specified parameter. However, such noise may still carry information about less frequently visited places. Such a possibility was explored by taking points classified in the previous step as noise and applying a less rigorous clustering via use of a significantly lower MinPts value of 500 (i.e. 40 min of combined presence). This procedure was repeated down to MinPts value of 150 (i.e. 10 min of combined presence).

Table 1
Identified clusters for one of the subjects

Cluster number	MinPts parameter	Description given by participant		
		Place	Times visited	Avg. duration of visit
1	5000	Workplace	Constantly	Hours
2	5000	Home	Constantly	Hours
3	500	Sports centre (gym)	Often	1 hour
4	150	Cafeteria	Once	Less than hour
5	150	City centre bus interchange	Several	Less than 30 minutes
6	5000	Friend's house (1)	Twice	Hours
7	150	Supermarket	Several	Less than 30 minutes
8	150	Bus stop	Several	Less than 15 minutes
9	150	Castle	Once	Hours
10	500	Club	Once	Hours
11	150	Church	Once	1 hour
12	500	Friend's house (2)	Once	Hours

Table 2
Repeated Bluetooth encounters. Number devices shown against number of appearances in separate scans

Appearances	Single $n = 1$	Occasional $1 < n \leqslant 10$	Regular $10 < n \leqslant 40$	Frequent $40 < n$
Devices	366	32	6	7

In order to verify the clusters, subjects were asked to identify places where clusters were found. In all cases, users were able to successfully identify discovered clusters as meaningful locations. Table 1 presents the results of cluster discovery and identification for one of the subjects. In general, clustering performed with greater values of MinPts identified frequently visited places with a large average occupation time such as home, workplace or gym. Lower values of MinPts enabled the discovery of less frequently visited places (e.g. one-off visits) as well as places that were visited frequently but for short periods of time (e.g. regular bus stops).

3.2. *Timeframes*

The above analysis takes no account of the time ordering of the data. Since the dates and times that places were visited could be highly relevant to both behaviour and changes in behaviour, then it is important to analyse the time course of locations. To investigate this, all clusters discovered with a MinPts parameter of 150 or greater were tagged by the participant as shown in Table 1. Figure 3 shows this data presented as a time series indicating the dates and times when the subject was in one of the recognised clusters. The figure also presents periods where existing readings were unclassified. Gaps in the plot represent times when the monitoring was essentially inactive, that is, when it had been turned off or no satellites were in view.

3.3. *Bluetooth encounters*

The dataset from which results have been presented also contained a record of Bluetooth encounters. The scans took place every ten minutes, whenever the application was running. During the trial, which lasted for 23 days, a total of 760 scans were performed that resulted in 411 different Bluetooth devices being detected. Perhaps unsurprisingly, as shown in Table 2, the vast majority of encounters were single

Table 3
Encounters with selected devices. Presence of the three most encountered devices in particular clusters

	Cluster 1 Work	Cluster 2 Home	Cluster 3 Gym	Cluster 4 Friend's home	Unclassified
Device 3 (Work PC)	346	0	0	0	10
Device 52 (Home PC)	0	169	0	0	24
Device 48 (Spouse's phone)	0	172	2	56	43

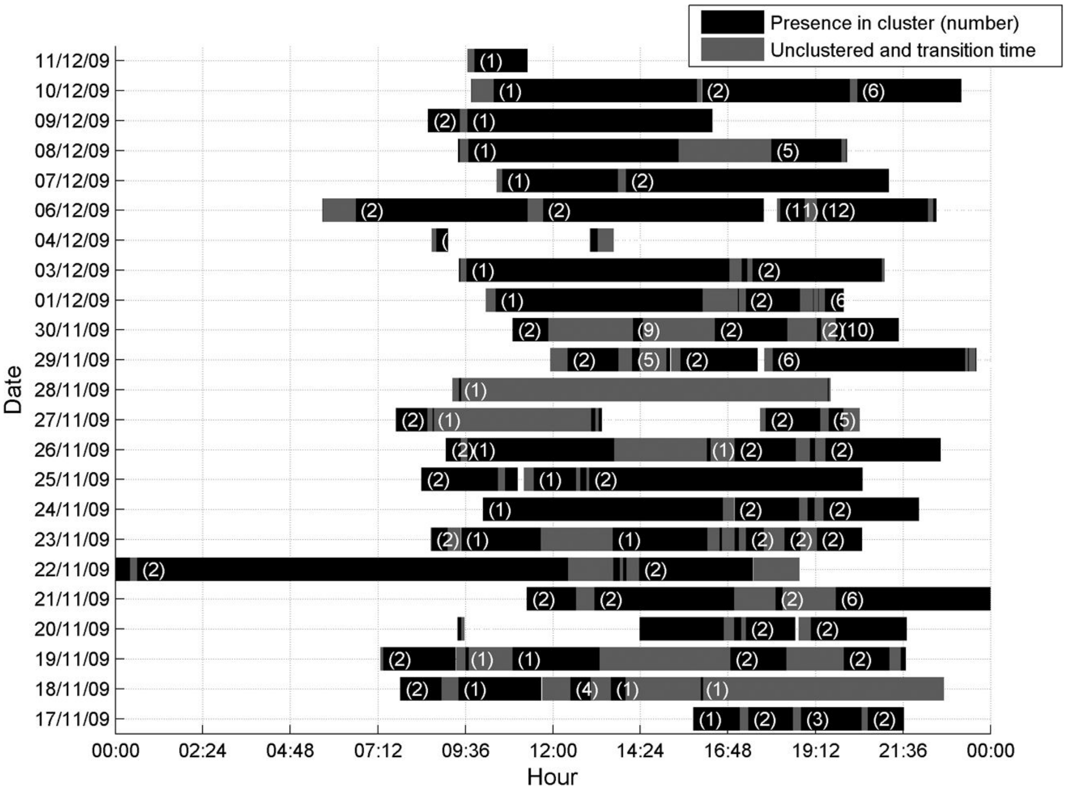

Fig. 3. Presence in clusters over time. The monitored person's presence in clusters over time. Each row corresponds to a day in the monitored period. Cluster numbers correspond to those presented in Table 1.

occurrences (366 out of entire 411). Of the remainder, 32 were 'occasional' encounters; 6 'regular' and 7 'frequent'. The most frequently encountered device, identified as a workplace PC station, appeared in almost 50% of all performed scans.

Encounters could often be linked to a particular cluster since the Bluetooth device may be a feature of the location's infrastructure, such as a Bluetooth dongle in a computer. When travelling with someone with a Bluetooth enabled phone, encounters in more than one cluster were obviously recorded. To further explore this, the most frequently encountered Bluetooth devices were identified and the number of times that they appeared in each cluster was calculated. Examples for three devices are shown in Table 3. Infrastructure related devices only appear in a single cluster whereas mobile devices can move from cluster to cluster as shown in Table 3. Clearly Bluetooth encounters can also provide location information, given certain assumptions and initialisation, whilst proximity to a mobile device along with location information can inform about social interaction [3].

4. Discussion and conclusions

4.1. Results overview

Applying proven clustering methods to GPS location data enables meaningful locations that a person has visited to be discovered. Repeated use of the method on the same dataset, but where previous clusters have been discarded enables locations that are significant on different levels, such as those constantly visited and those frequented occasionally to be discovered. This ability to prioritise discovered clusters may be important as bipolar manic episodes often result in the sufferers visiting new places such as restaurants or bars and spontaneously deciding to travel [6]. This would result in the number of detected 'infrequent' clusters growing and comparisons across time would show whether these were isolated visits or became meaningful locations over time.

During depressive episodes, the number of clusters appearing within a certain timescale would be likely to decrease significantly, due to the person's reduced social activity and diminished mobility.

Investigating the time-course of a person's presence in previously discovered meaningful locations may give additional information, since changes in behaviour may manifest themselves as variations in the timings of regular (from a spatial perspective) routines (e.g. a depressive patient may begin to stay at home during the day).

It has previously been suggested that monitoring Bluetooth encounters may provide a means of monitoring, relatively unobtrusively, social activity, a known indicator of mental health [5,10]. Such monitoring can give an insight into a person's daily activity to a degree that would be difficult and costly by other means.

What is more, applying knowledge about what particular devices are can enhance the process further. In the context of a behavioural disorder this might mean noting the unique Bluetooth addresses of friends and relatives and utilising this information for monitoring relationships, which can be highly affected during a bipolar episode.

4.2. Problems and data loss

The location data gathered during the trial was not complete and contained significant drop out in the data set. The main reason for this is the inability of the GPS receiver to work indoors. This did not actually influence the process of frequent location discovery as correct readings from such places were dense enough for clusters to be discovered. However, the density of points in the discovered clusters suggested that the time spent in those clusters was lower than in reality. Therefore, identifying occasionally visited indoor places may be problematic whereas less meaningful outdoor locations might be discovered (e.g. a long stay in a traffic jam).

Improvements could be made by utilising other technologies. Bluetooth encounters are one of the options as shown in this manuscript. However, Bluetooth infrastructure tied to a particular place is not that common and discovery available wireless networks (WLAN) to infer location using previous knowledge may be better, since such infrastructures are widely spread and discoverable in numerous places [14]. This would however, require a more advanced device than a mid-range phone as used here. Another possibility is to localise using cell information provided by the mobile network as shown in [15]. However, this method only approximates location with accuracy significantly lower than GPS (500m at best) and it is necessary to possess knowledge about cell locations.

Data loss due to loss of battery power was also the case. Frequent Bluetooth scans and maintaining communication with the GPS receiver causes the phone to discharge at a faster rate than usual. Generally, the phone needs to be recharged approximately every 24 hours whereas the GPS receiver needed to be charged approximately every 8 hours which may require changing typical usage habits. A possible solution for this problem would be to develop and apply algorithms that halt the GPS signal acquisition where location can be obtained by other means described above.

Another reason for data loss was because the users turned off the system in situations and places where they did not feel comfortable for monitoring to take place. This raises the issue of user compliance and privacy in location tracking although the pattern of cessation of monitoring may in itself contain valuable clinical data.

4.3. Privacy issues

Ethical issues concerning the proposed monitoring of a person's location is a complex problem that raises many questions [11]. Personalised ambient monitoring requires the processing of potentially sensitive data that may be shared with the patient's caregiver in order to provide individualised treatment. Therefore, efforts should be taken to maintain the patient's privacy to the highest possible extent. Therefore it is appropriate to process any sensitive data to retain relevant rather than precise information. For example exact position being mapped to a cluster number at the earliest opportunity. The methodology presented in this manuscript suggests how transition between anonymised clusters rather than specific locations is monitored.

4.4. Further work

The work presented here shows that currently available technology and techniques can be used to record and extract necessary geospatial information. Cluster analysis and assessing Bluetooth encounters has the potential to assist in the management of behavioural disorders as evidence shows that the detection of any changes in daily routine can be valuable [5]. Validation of the concepts presented here will be a step towards an effective warning and management system that could help improve a patient's quality of life and reduce the burden to the healthcare system [16]. To investigate the applicability and true value of the method described in this paper we now plan to embark upon a pilot trial with patients with bipolar disorder.

Acknowledgements

The authors would like to thank the EPSRC, UK for their funding of this project (Grant: EP/F003714/1), Jesse M Blum and James Amor for contributions to software design and conducting the study as well as David E Morris and Melissa Mather for valuable discussions on the manuscript text.

References

[1] G. Agamennoni, J. Nieto and E. Nebot, Mining GPS Data for Extracting Significant Places, *IEEE International Conference on Robotics and Automation, 2009. ICRA'09*, 2009, pp. 855–862.

[2] M. Bauer et al., Self-reporting software for bipolar disorder: Validation of ChronoRecord by patients with mania, *Psychiatry Research* **159** (2008), 359–366.

[3] N. Eagle, A. Pentland and D. Lazer, Inferring friendship network structure by using mobile phone data, *Proceedings of the National Academy of Sciences* **106** (2009), 15274.

[4] M. Ester, H. Kriegel, J. Sander and X. Xu, A density-based algorithm for discovering clusters in large spatial databases with noise, *Proc. Knowledge Discovery and Data Mining*, 1996, pp. 226–231.
[5] E. Frank, J.M. Gonzalez and A. Fagiolini, The importance of routine for preventing recurrence in bipolar disorder, *The American Journal of Psychiatry* **163** (2006), 981–985.
[6] S.N. Ghaemi, Mood Disorders: A Practical Guide, London, Blackwell, 2008.
[7] B. Hofmann-Wellenhof, H. Lichtenegger and J. Collins, *Global Positioning System: Theory and Practice*, Wien, Springer 2001.
[8] C. James et al., Personalised Ambient Monitoring (PAM) of the mentally ill, in *Proc. 4th European Conference of the International Federation for Medical and Biological Engineering ECIFMBE*, Antwerp, 2008, pp. 1010–1013.
[9] V. Kostakos, E. O Neill and A. Shahi, Building common ground for face to face interactions by sharing mobile device context, *Lecture Notes in Computer Science* **3987** (2006), 222.
[10] S. Malkoff-Schwartz et al., Social rhythm disruption and stressful life events in the onset of bipolar and unipolar episodes, *Psychological Medicine* **30** (2000), 1005–1016.
[11] K. Michael, A. McNamee and M.G. Michael, The Emerging Ethics of humancentric GPS Tracking and Monitoring, in *Proc. International Conference on Mobile Business*, Copenhagen, 2006.
[12] R. Morriss, The early warning symptom intervention for patients with bipolar affective disorder, *Advances in Psychiatric Treatment* **10** (2004), 18.
[13] P. Nurmi and J. Koolwaaij, Identifying meaningful locations, *Proc. 3rd Annual International Conference on Mobile and Ubiquitous Computing* (*Mobiquitous*), San Jose, CA, USA, Citeseer, 2006.
[14] J. Rekimoto, T. Miyaki and T. Ishizawa, LifeTag: WiFi-based continuous location logging for life pattern analysis, *Lecture Notes in Computer Science* **4718** (2007), p. 35.
[15] A. Varshavsky, M. Chen, E. de Lara, J. Froehlich, D. Haehnel, J. Hightower, A. LaMarca, F. Potter, T. Sohn, K. Tang and others, Are GSM phones THE solution for localization, *Proceedings of the Seventh IEEE Workshop on Mobile Computing Systems & Applications*, Citeseer, 2006, pp. 20–28.
[16] H. Wittchen and F. Jacobi, Size and burden of mental disorders in Europe – a critical review and appraisal of 27 studies, *European Neuropsychopharmacology* **15** (2005), 357–376.
[17] C. Zhou, S. Shekhar and L. Terveen, Discovering personal paths from sparse gps traces, 1st International Workshop on Data Mining in conjunction with 8th Joint Conference on Information Sciences July, Citeseer, 2005, pp. 3–6.