

Shakra: Tracking and Sharing Daily Activity Levels with Unaugmented Mobile Phones

Ian Anderson · Julie Maitland · Scott Sherwood ·
Louise Barkhuus · Matthew Chalmers · Malcolm Hall ·
Barry Brown · Henk Muller

Published online: 3 August 2007
© Springer Science + Business Media, LLC 2007

Abstract This paper explores the potential for use of an unaugmented commodity technology—the mobile phone—as a health promotion tool. We describe a prototype application that tracks the daily exercise activities of people, using an Artificial Neural Network (ANN) to analyse GSM cell signal strength and visibility to estimate a user’s movement. In a short-term study of the prototype that shared activity information amongst groups of friends, we found that awareness encouraged reflection on, and increased motivation for, daily activity. The study raised concerns regarding the reliability of ANN-facilitated activ-

ity detection in the ‘real world’. We describe some of the details of the pilot study and introduce a promising new approach to activity detection that has been developed in response to some of the issues raised by the pilot study, involving Hidden Markov Models (HMM), task modelling and unsupervised calibration. We conclude with our intended plans to develop the system further in order to carry out a longer-term clinical trial.

Keywords activity recognition · context aware · daily activity levels · mobile phones

I. Anderson (✉) · H. Muller
Department of Computer Science, University of Bristol,
BS8 1UB Bristol, UK
e-mail: anderson@cs.bris.ac.uk

H. Muller
e-mail: henkm@cs.bris.ac.uk

J. Maitland · S. Sherwood · L. Barkhuus ·
M. Chalmers · M. Hall · B. Brown
Department of Computer Science, University of Glasgow,
G12 8QQ Glasgow, UK

J. Maitland
e-mail: jules@dcs.gla.ac.uk

S. Sherwood
e-mail: sherwood@dcs.gla.ac.uk

L. Barkhuus
e-mail: barkhuus@dcs.gla.ac.uk

M. Chalmers
e-mail: matthew@dcs.gla.ac.uk

M. Hall
e-mail: mh@dcs.gla.ac.uk

B. Brown
e-mail: barry@dcs.gla.ac.uk

1 Introduction

The decreasing levels of daily activity undertaken by the general public form an ongoing challenge for those involved in public health, and are of concern to both primary and secondary healthcare. The benefits of physical activity are well documented and widely acknowledged, and yet the World Health Organisation states that 60% of the worldwide population are not active enough to profit from these benefits [22]. Pervasive and ubiquitous computing technologies are well-suited for use within the healthcare industry and have the potential to be far-reaching and effective. This paper presents a prototype application that runs on arguably the most pervasive computing technology of all, the mobile phone. By detecting patterns in signal strength fluctuation and changes in the visibility of GSM cells, the application uses an ANN to infer whether the carrier of the phone is sitting still, walking, or travelling in a car. This information is then used to calculate the carrier’s daily activity level, which can then be shared with and compared to the activity levels of others.

Rather than being driven by experimental hypothesis or outcomes, we took a more exploratory approach. A short-term pilot study gauged information regarding usability, user response and attitudes toward the prototype, and will inform the future refinement of the system so that it is suitable for longer-term clinical trials.

Since the pilot study took place, further work has been done to improve the reliability of the system's activity detection which, although performing adequately for the purposes of the trial, seemed to have difficulties when processing GSM data from disparate environments: the patterns of GSM behaviour obviously differed for similar physical activity when situated in rural as opposed to urban environments. To overcome this issue, automated unsupervised calibration based on HMM enables the dynamic learning of mappings between GSM behaviour and physical activity. This method is explained in detail and compared to our original approach in Section 7.

1.1 Background

Augmentation of traditional exercise technologies and practices with more pervasive and ubiquitous computing is becoming an established area in both research and commercial arenas. Gym-based equipment such as treadmills and rowing machines has been complemented with virtual reality environments to motivate and stimulate users during their workout (<http://www.fpgamerunner.com>). Some of this equipment also facilitates the downloading of workout data onto an individual's PDA, so that personal workout programs may be monitored and adapted (<http://www.acrocat.com>). Positioning technology is now also being used to extend such services beyond the gym environment to benefit walkers, cyclists and road-runners. These technologies assist individuals who have already taken steps to get fit or remain healthy. However, there is relatively little in the way of assistive or motivational technologies that are aimed at the more sedentary adult/child.

Self-monitoring is a well established behavioural change technique [13]. Pedometers are illustrative of this technique and are a widely used fitness-related technology that does not demand a vested interest in health in order to be used. Although their accuracy may be volatile, they have been found to motivate individuals taking early steps towards a more active lifestyle [19]. As pervasive technologies advance, so does the ability to detect and monitor the physiology and physical activity levels of an individual or community to an increasingly fine granularity. A multi-modal sensor board can now distinguish between eight physical activities [11], and commercially available technology can be worn on the body to monitor blood pressure, heart rate, and stress levels. These technologies are

unarguably useful, but their specialist nature may prove to be a barrier to widespread adoption and utilisation.

The recommended level of activity for an adult is at least 30 min of moderate activity, five times a week. Although prolonged periods of activity are most advantageous, the daily amount of 30 min can be accumulated throughout the day in shorter periods of 10 min or more [6]. Most adults who do not currently reach this level of activity may be able to achieve this target by making small changes to their everyday routine. By capturing and acknowledging everyday activity in an accessible and noninvasive manner, and facilitating the sharing and comparison of that information between peers, we hope that awareness will be raised in such a way that it motivates users to become more active on a day-to-day basis. Shakra is our first prototype of such a system, and it runs on an unmodified mobile phone. Although not everybody owns a mobile phone, it is the most uniformly adopted technology throughout all social classes [8], and so this platform hopefully overcomes the aforementioned barrier to adoption and, therefore, effect.

The following section reviews conceptual approaches to behavioural change in physical activity, alongside current technical approaches to fitness tracking and motivation. The resultant design and implementation detail of Shakra follows, before the pilot study is presented and discussed. Following the discussion of our study, refinements and implications for further development are presented, which are aimed at improving the Shakra system for a larger clinical trial. These are also presented as guidelines for future development of similar health related systems.

2 The problems of motivating exercise

Numerous studies show how just a minimal amount of daily activity can increase general health [15, 18]. Having a lifestyle that promotes regular exercise seems to be a challenge in the western world, since our daily lives are busy and many of us draw upon transportation systems, such as cars, trains and buses. Studies suggest that around 70% of the UK population fails to meet minimum recommendations for physical activity [1].

In view of the aforementioned recommendations, many approaches to increasing fitness propose an increase in *moderate* activity, such as brisk walking, in order to improve people's health [7]. Moderate activity is generally defined as when a person's heartbeat is increased to 55–69% of maximum heart rate, which for most people would occur when walking at about 4 mi/h. One important factor for consideration is that many people have difficulties making sure that their activity is in fact moderate and not just light, i.e. that they are achieving the health benefits

stated above [14]. It is therefore important for individuals to not only be aware of their overall amount of exercise but also its intensity.

2.1 Tracking and motivating fitness and moderate activity

Many technical methods have been developed to measure fitness and physical activity. In contrast to physiology-oriented systems such as Qualcomm's Cardionet (<http://www.qualcomm.com>) and low-cost pulse oximeters that enable users to monitor their blood pressure, heart rate and blood oxygen saturation, more behaviour-oriented systems focus on raising awareness of their users' behavioural state. One common device is the pedometer, a small device that measures each stride the wearer takes. One recent report indicates that just the presence of the pedometer can motivate people to be more active [19]; another study showed that sharing daily step count within a small group of friends was more satisfying and motivating compared to a control group who measured but did not share their information [5].

One of the most advanced commercial technologies in this area is the BodyBugg (<http://www.bodybugg.com>), also known as SenseWear. The BodyBugg measures an array of values such as relative body temperature, step count and acceleration, in order to estimate how many calories the wearer is burning. It has been shown to work reliably in controlled tests for measuring calories burned, with an accuracy of 89–98%, however it is limited in its determination of the actual context of the wearer [11]. Also, it has to be worn on the upper arm for 24 h a day; it can therefore easily disrupt sleeping and collide with everyday clothing—a particular disadvantage among women who often wear tighter or lighter clothes.

A less direct means to motivate activity is taking part in mobile games. Most mobile games involve infrequent play over a relatively short period, with limited health benefits, but some games such as Mogi Mogi (<http://www.mogimogi.com>) and Feeding Yoshi [4] take place over a longer period of time and are more 'interwoven into everyday life'. A study of Mogi Mogi showed that players would frequently take detours from their normal routes, and that "many alight at an unusual metro station on the way home if they notice an object on their mobile screen, even if this means walking much further to get home. Many players also said they went out at night because the mobile screen had indicated objects in the vicinity" [12]. Similarly, Bell et al. report that players adjusted their everyday routines of work and travel so as to spend more time playing the game, often walking a good deal more than they would do normally. A disadvantage of this approach is its relative lack of clarity or precision about the exercise undertaken. While players increase their activity as part of playing the game, this is not directly

connected to or encouraged by the game—instead it is a useful but indirect benefit of the game.

2.2 Theories and studies of change in activity

Numerous studies have explored how to motivate people in increasing their activity level, and there are two well-cited theoretical approaches: the Transtheoretical Model, where behaviour change is described as a multistage process [21] and Social Cognitive Theory, based on the individual's outcome expectancy and self-efficacy [18].

The *Transtheoretical Model* is one of the more common theories referred to in the health literature. It focuses on the individual stages people go through with regard to physical exercise regimes, such as pre-contemplation, contemplation, preparation, action and maintenance. Although it is possible to determine people's individual stage at a given time with a standard questionnaire, the theory does not account for individuals' different *levels* of exercise and it does not address the possibility for individuals to skip between the stages. One critique has also been that it is focused on attitude rather than behaviour, although in observational terms both seem to be significant. For example, it has been pointed out that the difference between the stages of pre-contemplation and contemplation only refers to a change in attitude rather than actual change in physical activity. Moreover, recent research points to the theory's weakness in showing long-term changes [1].

The *Social Cognitive Theory* focuses on increasing the individual's self-efficacy by different means, in relation to keeping fit, leaning on studies that show how intrinsic motivation (enjoyment, feeling good about the exercise) rather than extrinsic motivation (external pressures or immediate rewards) increase the likelihood that the person will stick to a routine [14]. Examples of intervention using this approach include giving health advice over the phone, either by health professionals or via an automated service, and through an Internet service [10]. Studies showed that human interaction for example was successful in promoting increased physical activity among middle-aged and elderly when compared against a control group [9].

Other research has addressed social aspects of sharing information about activity and found that exercising together can also motivate individuals to do *more* activity; people increase their activity level as they engage in the light competition [20]. Similarly, when people receive tailored information that is personally relevant, it is more likely to stimulate change, adding to people's self-efficacy and outcome expectancy [21]. It is evident that intrinsic motivation is influenced but not determined by wider social interaction.

Behavioural change is difficult to promote, and many researchers point to the combinations of internal and

external influences that are complicated to trace, target and categorise in individual cases. One critique that has been made of the physical activity literature, for example, is that it does not separate between individual environmental values (such as age, social class, health status) and social environmental values (such as family, school/work and community) [7]. The social cognitive theory addresses aspects of community, contrasting to the Transtheoretical model although it focuses on internal values as main motivator to increase individuals' level of exercise. Interestingly enough, social factors such as poverty and neighbourhood have been found to highly influence people's level of exercise [7]. Again, such categorisations abstract over individual cases, but it is reasonable to conclude that an individual's exercise is often affected by interactions with his or her surrounding group. In our work, we therefore focus on social and communal aspects of exercise; the light pressure from the surrounding community is a great motivational factor not to be underestimated in relation to intrinsic motivational factors. Also, rather than taking a broad survey and relying on social categories such as class, in our evaluation we focus on the details of particular individuals' experience. Based on our understanding of related theories and studies, our system design is directed towards a long-term goal of achieving greater public health. We assume that a member of the general public is likely to make only minor behavioural changes, and that this will be based on individual awareness as well as social interaction.

3 The Shakra prototype

Our overall aim is to design and implement a system that will help to motivate adults who do not currently achieve the minimum recommended daily activity level, and who can benefit from a raised awareness of their current levels of activity: a system that can track and categorise an individual's daily activity into accumulative time spent in inactivity, light, moderate, and vigorous activity. In acknowledgement of the influence that social networks can have on the actions of an individual, the system should facilitate the sharing and comparison of data between peers. In order to evaluate user response to such a system and general usability, a basic prototype was created that determined whether a user was active or inactive, accumulated daily totals, and allowed the sharing and comparison of the daily totals.

A client-server architecture was used in Shakra (see Fig. 1). The client was built for Windows smartphones running Windows Mobile 5 (WM5). In the system trial the particular phones used were i-mate sp5s, which had SIM cards enabling them to connect to a commercial mobile phone network for data transfer. The client was implemented in C# using the .NET compact framework. A web

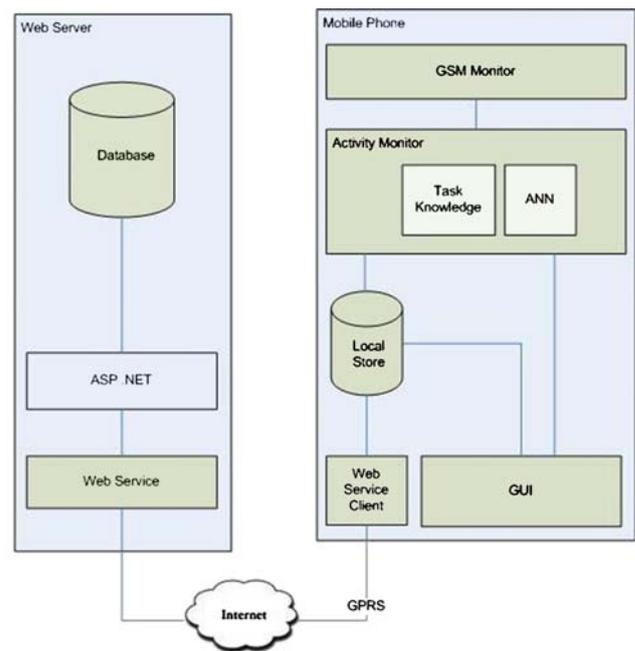


Figure 1 System architecture

service was used so that clients could upload their own and download others' data. The web service ran on a server running Windows Server 2003 and the data uploaded was stored in a MySQL database.

Our key design goal for Shakra was that it could be carried around in a nonintrusive manner, requiring little or no extra equipment for users. Minimal user intervention is required in order for it to function effectively; the system tracks the activity of the user without direct manual input. The application tracks users' general level of activity, showing the current mobility state: no movement ('stationary'), moderate activity ('walking') and travelling in a car, bus or train (collectively labelled here as 'driving'). The moderate activity is then used to display a 'minutes of activity per day', with a historical view supporting comparison of activity across the previous week. This supports a user monitoring his or her own activity and exercise levels, with the exception that stationary exercise (such as working out at a gym) is not tracked.

When running the application for the first time, the user is prompted to provide a name, used to identify himself or herself within the system and to other users. The application records up to seven visible GSM cells and their signal strengths, once per second. The current activity of the user is then classified every 30 s by the application's neural network, as described in more detail below. Using a web service, each phone uploads the recorded activity of the user via GPRS and stores it in a MySQL database, while simultaneously downloading information about other participants for later review. The system updates this shared information automatically every hour. If a user does not

want to wait for an update, he or she can manually synchronise via the *Sync* menu option.

Users specify in advance the peers they wish to share results with, but at any time they can change this list. Figure 2a shows the *Compare Daily Activity* screen that users can view to assess their performance in relation to their peers. For a week's overview of their own activity, users may use the *Week's Activity* screen shown in Fig. 2b. In order to provide real time feedback to the user, an animated representation of the user's current mode of activity runs continuously on the main screen of the application. This is shown in Fig. 2c and d.

3.1 Sensing activity

The current activity of the user is inferred using patterns of fluctuation in GSM signal strength and changes to the IDs of detected cells. This method has been demonstrated as a reliable and unobtrusive way of sensing current activity [2], and has the advantage over the more traditional approach of using an accelerometer in that it does not require additional sensor hardware as in Sensay [17] and the multimodal sensor board of [11]. Similarly, while the processing of physiological and biometric data could complement our approach, the benefits of encapsulating the system within a mobile phone would be lost. An alternative approach would be to utilise the positioning information available from some mobile phone networks, however this approach frequently involves prohibitive cost, as well as depending upon much of the same technology as our client based monitoring.

Rather like a traditional accelerometer, the levels of signal strength fluctuation change when a mobile phone is moved. For example, Fig. 3 shows the total signal strength fluctuation across all monitored cells during successive 30-s time periods whilst walking, remaining still and travelling

in a motor car. The figure illustrates that it is relatively easy to distinguish between moving and remaining stationary, but at times, the pattern of fluctuation whilst walking will match that of driving and vice versa. This is due to the stop–start nature of both walking and travelling in a motor car in urban areas. When driving, a greater geographical distance will typically be covered over a given time period when compared to that of running or walking. As such it is possible to use the rate of change of neighbouring cells to infer travel by car.

To classify these patterns we use an artificial neural network. The network inputs are the sum of signal strength fluctuation across all monitored cells, and the number of distinct cells monitored over a given time interval. The network consists of a single layer of eight hidden neurons; weights are learnt using back propagation. The network outputs the currently sensed activity for the given input values. The network is trained by repeatedly presenting data collected during each method of movement.

The current activity of the user is conditionally dependent upon their previous activity. In order to provide instant feedback to the user interface, the neural network deliberately does not model this behaviour. Instead, when determining if any additional minutes have been earned, we apply task knowledge based upon the output from the neural network over the previous two and a half minutes. This enables noise to be filtered out and a more accurate representation of the users' activities achieved. For example, periods of low signal strength fluctuation such as stopping at traffic lights whilst driving can be ignored when placed between periods of high fluctuation where many distinct neighbouring cells were monitored. It could be argued that activity would be more accurately inferred if a longer rolling filter had been applied to the GSM data. Introducing longer filters would have increased the likelihood of active minutes 'disappearing' from the users'

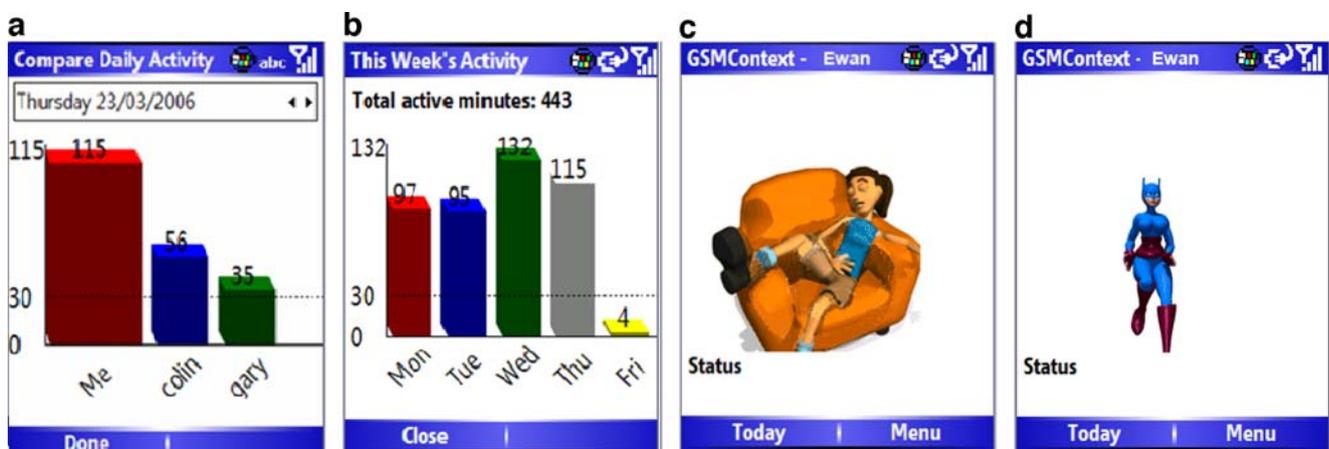
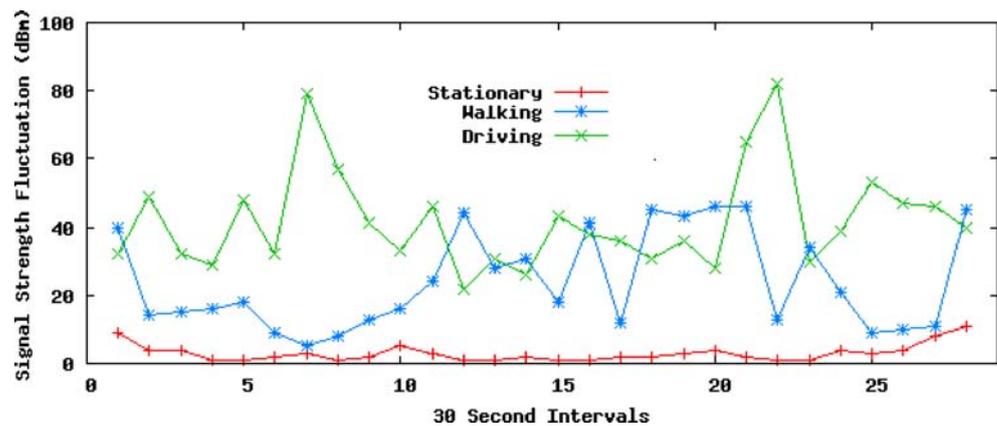


Figure 2 The phone interface. Images a and b show screens for examining relative and individual activity levels: compare Daily Activity and This Week's Activity Images. c and d show two of the screens showing the estimated current activity level: Stationary and Walking

Figure 3 Distinguishable patterns of GSM signal strength fluctuation over successive 30 s samples are used in identifying the activity levels Stationary, Walking, and Driving



activity totals. A decision was made that, for the purpose of this study, priority would be given to user experience, with the intention that this trade-off would be addressed in future work.

4 The user study

The Shakra application was evaluated with three groups, to detail its use, to determine whether it increased users' awareness of their activity level and if this could potentially motivate them to be more active, and to derive implications for future work. Naturally, a longitudinal clinical study lasting months or years would be needed to rigorously assess long-term changes in users' behaviour and health, but our one week trial served as a pilot evaluation of a potentially powerful activity promoting application. The focus was on the users' experiences with both the activity tracking and the sharing feature; it was important to find if sharing information was good for increasing awareness and motivating a more active lifestyle.

Before the trial, a base neural network had been constructed by using GSM data collected by the development team while sitting still, walking and driving. In order to determine whether or not further personalisation of the network was required for each of the trial participants, the system was given to each participant for a 2-day training period. During this period, the participants were asked to record whenever their activity mode changed. Functionally, this was a simple task supported in the application's main interface that users learned to do quickly. For the training days, we asked the participants to take the phones with them as they went about a normal day's activity. This trained the system for the areas in which they usually go throughout the course of a day.

Following the initial system-training period, the data collected by the trial participants were analysed. We found that only minor changes to the previously trained neural network were required by three of the nine volunteers. This

was due to them living and working in urban areas that exhibited different levels of signal fluctuation to those where the initial training data had been collected by the research team.

4.1 Method

Overall, the trial took place over 10 days. The participants initially filled in a simple activity diary for 3 days, to determine their present levels of activity and to compare activity to the week of using the application. Immediately after, they trained the system for 2 days and then finally used the system for a 5-day working week, filling in a diary describing their use of the system and whereabouts for each day. We kept in touch with the participants by phoning them once during the week, and sending text messages in the few cases where it looked like the phone was not uploading properly. At the end of the study, each participant was interviewed individually to expand on the use and reflect on the experiences with and opinion of Shakra.

The participants were recruited as groups of friends and/or coworkers who had daily interaction with each other and would enjoy sharing their exercise information. Although the target users for the final system are inactive people, it is unrealistic to expect that only inactive people will use the system. This is especially true when the system is aimed towards peer groups who will naturally include individuals of differing levels of activity. We therefore aimed to study the use of the system among a diverse set of people, and the nine participants varied in the degree of their normal activity. Two were highly active, with purposeful exercise at least 3 days a week, four were moderately active people, working out one to two times a week, and three were fairly inactive, walking but not doing any purposeful exercise¹. Table 1 provides an overview of the three groups.

¹ Naturally this is a very broad characterisation from the participant's own statements and diary reports. It is not necessarily a true reflection of their level of health or level of fitness.

Table 1 Participants in each group

Characteristics	Group 1	Group 2	Group 3
<i>N</i>	2	3	4
Age range	52–54	28–30	19–37
Sex	Female/male	Male	Female
Activity level	Fairly inactive	Two moderately and one highly active	One inactive, two moderately active and one highly active
Occupation	Teacher and administrator	Technical administrators	Manager, administrative staff and student

After the study, the system logs were analysed. First of all, the activity times were compared to the self-reported diaries and the interviews, to make sure there was a fair level of accuracy in measuring activity. Secondly, the logs were scrutinised to see how participants used the application, how often they compared their activity to others', and how often they looked at their weekly chart. The interviews were transcribed immediately and the parts were categorised according to major topics and themes. They were used to elaborate on the diary, such as precise times of commute, actual transport methods, and more detailed experiences and impressions of the application during the week. In the next section, we report the results in relation to three topics, one relating to precision or reliability of the application's measurements, a second looking at users' individual experiences, and the third exploring the participants' experiences of information sharing.

5 Reliability of Shakra in the real world

Although previous tests had shown highly accurate determination of activity [2], the real test of the application would be using it in an uncontrolled environment among many different people. We did not expect to get as high accuracy, because of the unstructured and diverse behaviour of people leading their everyday lives. Overall, the application showed very good determination of activity and the participants found it very useful as a tool for measuring their activities. After analysing the diaries and annotating them with information gained through interviews, we compared each day of each participant with a log-generated activity timeline. It was easy to see participants commute to work, break for lunch, and commute back from work; two examples, with diary annotations, are shown in Fig. 4. A rough analysis was done to determine the rate of correct labelling of activity. We chose three sample days for two

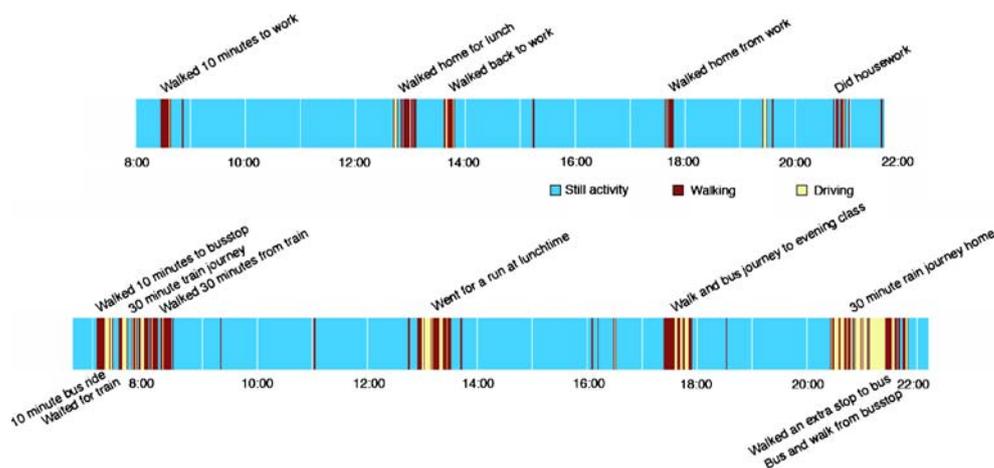
different participants because their diary entries for those days were particularly comprehensive, i.e. 6 days in total. From the unfiltered data we analysed short stretches of 60 to 90 min with varied activity; this was done to refrain from considering the long hours of inactivity, which occurred during their work day where participants were mostly sitting at their desk. Including this would have given unrealistically optimistic numbers. Results showed a minimum of 70% accuracy, during users' commute when fluctuations are highest. The misinterpretations often occurred during changes between different methods of transportation such as getting off a bus or a train. However since there would often be a delay both before and after transportation, the misinterpretations would cancel each other out, correcting the accumulated minutes of exercise. One more problematic finding was that running would occasionally register as driving. During one participant's 45-min lunch run, 15 of the minutes were registered as driving. For another participant with a long commute, for example, it meant that he gained a maximum of seven active minutes each day due to error. This was the maximum error we found from looking at participants' commutes.

Some of the diary entries assisted in showing when still or walking activity was misidentified. For example, one woman from group 3 explained that she went on a walk for 30 min, but had only increased her overall activity count by 22 min when she returned. It should be noted that this particular participant lived in the countryside where we knew that the neural network would be less accurate. Similarly, a male participant reported that his 10 min walk to work sometimes only gave him 7 to 8 min of activity. This may in part be attributable to a lag in activity determination, as well as the participants stopping at road crossings, etc. Since the application is aimed towards increasing awareness rather than measuring physical exercise precisely, and offered useably accurate overall measures, we suggest that the small moment-by-moment lags and jitters in classification were not problematic for the purpose of the trial. Real-world reliability is, however, essential to enabling a broader range of applications, especially those involving moment-by-moment tracking and display. Since the pilot study took place, alternative methods of activity detection have been evaluated and Hidden Markov Models have been found to be a promising substitute to the current Artificial Neural Network implementation. Our findings are presented in detail in Section 7.

6 User experience

The participants all took the phones with them every day, carrying the phones around with them wherever they went

Figure 4 Example timelines of activity for two participant's days with colour showing the activity level and text showing the participants' diary annotations



for the vast majority of the day. The application was found to be both reliable and stable overall, and everyone found it easy to use. Where group 2 had the chance to use it during most of their working day, and therefore checked it and compared extensively (between 11 and 34 times a day), the other groups had busy days where they would mostly check their numbers and compare in the evening, therefore checking fewer times (between 1 and 20 times).

Participants reported that the application was fun to use and gave them good—and sometimes surprising—awareness of their activity level. Two participants (from groups 2 and 3) reported it to be highly ‘addictive’, in particular the sharing aspect. Another participant repeatedly explained how it made him see how ‘lazy’ he was. Although only four of the nine participants reported doing more activity than usual in the interviews (and attributed it to the application’s sharing functionality as well as more general competitiveness), the diaries show that the other participants were also more active compared to the initial three day ‘base’ diary. The short-term nature of this pilot study does not allow for observation or inferences to be made about the initial novelty value of the system. As we discuss further in the Conclusion and Future Research section, a longitudinal clinical trial will determine long-term use and effects.

6.1 Individual use and motivation

The participants described how they would enjoy checking how much walking and running activity they did during the day. Most of them checked their own minutes regularly and were astonished how they gained minutes during busy days. One woman from group 3 was surprised that she had accumulated 177 min one day, but when looking back though the diary, she realised that she had been busy commuting between two different work places (which involved walking to and from a bus and a ferry), as well as walking her dogs in the morning and evening. We were

able to detect most of her activities in the data log, except for some of her transport that had a few small gaps of 30 s walking when she was in fact driving. This error, however, did not add more than 7 min of walking to the whole day. This participant was busy and already highly active, and did not feel the application had made her change her activity level during the study.

One participant from group 2 on the other hand, was very active that week in particular, and attributed this to the application. He explains how he increased his activity that week:

[I]t probably encouraged me to go running Monday, Wednesday and Friday, because I always have the intention of going running at the beginning of the week. [...] and I sort of set out Monday, okay right, I will take my stuff and I will go, you know, just Monday, Wednesday and Friday. [It also encouraged me to] just walk a couple of extra bus stops [...]

He was very keen on increasing his activity level, and had tried to get into running 3 days a week for a while, without complete success. The weather had sometimes deterred him before, but with Shakra, he went out every planned day despite it being very rainy two of those days.

Although the participants seemed to be motivated from just the awareness of their activity, the effect was not unanticipated; often merely the knowledge that others can detect one’s activity (either from a fill in diary or a tracking system) makes one more active. However, it was important to explore whether the use affected users’ awareness and attitude towards moderate exercise. Behavioural change is a slow and often long-term process, but the necessary first steps have been taken here, in that awareness and motivation increased. Other issues affect motivation and awareness in return; therefore it should be related to social factors such as competition and collaboration—as the next section discusses.

6.2 Shared experience

The groups did not only enjoy the increased awareness of their individual activity levels, they also enjoyed competing among themselves. Group 2 were quite determined in their competition, in particular one participant who would spend much of his working day walking around taking calls on his wireless headset, much more than he usually did. Another of his group's members explains:

...[W]e would be sitting in calls and he would be walking by [showing the phone to us]. Maybe there was a meeting round that side of the building (pointing), he would walk all around the building to get there (the building is doughnut shaped) ... Me and Colin would sort of check more often to see. Ewan just rubbed it in front of our noses, how far he went.

This group enjoyed the competition despite a very different number of accumulated active minutes as Fig. 5 shows. Since the 'overachiever' described above had a wireless headset and was not confined to his desk, he could work while walking around—or walk while working. The other two group members were more confined to their desks during the day and only reached about half of his minutes every day. Where the first of these two said that he realised how 'lazy' he was. The participant second explained that he did not care that much, since he worked out at the gym about three times a week. He was quite content with his activity level, and did not see his 10-min walk to and from work as 'exercise'. In this case there was more concern from the less active of the two, who was in the category that the application is most focused on, although he was constrained in changing this awareness into greater activity—at least during the trial.

Group 3 also started competing, with two women particularly competitive with each other. One wanted to beat her very active friend. For example, one evening when she came back from a run with 112 min, she saw her friend had 177 min of activity. In an attempt to catch her friend up, she asked her neighbour if she could take the latter's

dogs for a walk. She therefore managed to get 137 min—not quite enough to beat her friend, but a respectable amount of exercise to say the least.

Group 1 did not compete much, but they did enjoy the fact that they could see each other's activity when they were apart. The oldest of the study participants and also a married couple, they mostly used the system to keep an eye on their own activity levels.

6.3 Sharing the fun

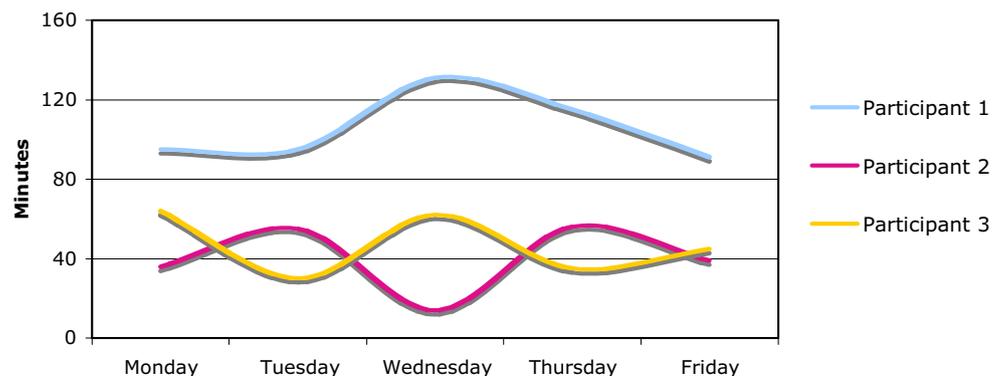
One distinct difference between our application and mobile games is the designed purpose of promoting exercise versus promoting play. However, we found that the difference in use is not necessarily so distinct. Where other games have been shown to promote exercise, we found that playfulness can be a side effect of health-focused applications.

Participants had fun competing as described above and they did not only use it for teasing each other and as a conversation topic: some of them saw it as a game. One of the participants commented that his buddy "wanted to win so much. Before we could even get it to a certain level, he was flying", he said. 'Walking around' with the sole purpose of gaining active minutes was common behaviour among some participants, which not only shows their competitiveness but also how they wanted to 'play' the system. In essence, the application has game-like characteristics for those who like to play: the winner is the person who accumulates the most activity in a day.

7 Post trial enhancements

Although generally not perceived as a problem during the pilot study, issues of accuracy were raised by the apparent lags and jitters in the system's classification of the user's current activity. Sudden spikes and troughs in GSM cell strength and visibility contribute to the level of difficulty associated with GSM cell-based activity inference. In Section 5 we discussed the need to analyse previous mea-

Figure 5 Group 2's accumulated minutes per day



surements alongside those currently detected in order to filter the noise associated with the GSM data. During the trial this was implemented as a simple smoothing algorithm using the output from the neural network. In this section we process data obtained during the trial using a Hidden Markov Model (HMM) and present an unsupervised approach to calibration. We demonstrate how a HMM based approach to inferring activity is more appropriate than an ANN because the HMM can directly model the task knowledge. We conclude this section by comparing results from the HMM and the ANN obtained by processing the same trial data.

7.1 Hidden markov model (HMM)

The problem that we are trying to solve is that we wish to infer activity of the cell phone carrier from observations of the GSM data. The GSM data provides with an indication of the activity, but this needs to be smoothed out by knowledge of “normal” behaviour. For example, it is usual for a person to drive for a prolonged period of time, and then to walk; it is unusual for a person to switch between driving and sitting frequently. We can model this activity using a Hidden Markov Model (HMM). A HMM λ is defined as follows:

$$\lambda = (A, B, \pi)$$

A is the transition matrix representing the probabilities of moving from one state (activity) to another. In the context of HMMs, the activity of a cell phone carrier is referred to by the term *state*, that is, the hidden nonobservable state. Therefore for the rest of this section we use the term state to refer to carrier's activity. B is the observation matrix representing the probability of being in a state given an observation and π is the initial probability distribution. S represents the set of states that the carrier can be in (the *state alphabet*), in our case:

$$S = (\text{still}, \text{walking}, \text{driving})$$

V is the set of discrete observations. It comprises n elements (v_1, v_2, \dots, v_n) . In our case, we map measurements of the signal strength fluctuation and the cell fluctuation onto a set of 15 discrete observations. This process is described in detail in the following section. During operation of the HMM we will have a sequence of observations which will lead to a sequence of states. There are t observations O , and t matching inferred states Q :

$$O = (o_1, o_2, \dots, o_t) \quad Q = (q_1, q_2, \dots, q_t)$$

The strength of a Hidden Markov Model is that it uses knowledge of previous states in order to predict the most probable current state. In our case, we remember five prior states, and this is represented in the transition matrix A .

Hence, the probability of q_t depends on states $(q_{t-5}, q_{t-4}, \dots, q_{t-1})$:

$$P(q_t/q_{1-t}) = P(q_t/q_{t-5}, q_{t-4}, \dots, q_{t-1})$$

The matrix A captures these probabilities: it contains the probability of transitioning to state j given the previous five states of activity in a sequence, $q_{t-5}, q_{t-4}, \dots, q_{t-1}$, that is:

$$a_{ij} = P(q_t = s_i/q_{t-5}, q_{t-4}, \dots, q_{t-1})$$

where q_t is the future state in a sequence. The observation matrix B contains the probabilities of an observation k being produced whilst currently in state j :

$$b_{jk} = (v_k/s_j)$$

In order to infer the most likely state sequence given a sequence of observations we can use the Viterbi algorithm [16, 23].

7.2 Unsupervised calibration

In this section we present a method for unsupervised calibration using the HMM described in the previous section. We use the Baum-Welch method to learn A , B , and π [3]. By presenting the Baum-Welch algorithm with a sequence of observations it will populate A , B , and π . It will not however help us consistently map signal strength and cell fluctuations to the observation alphabet. If this mapping is not consistent across environments then the inferred hidden state may imply a different meaning such as walking when the user is actually driving. To avoid an arduous data collection and calibration procedure we use an *automated, unsupervised process* to learn the mapping between cell and signal strength fluctuation and the observation alphabet mapping. In this section we describe that process.

Each of the activities we aim to distinguish between produces a different pattern of signal and cell fluctuation. The pattern of fluctuation depends upon the environment. By identifying the mean fluctuation values for activities in specific environments we can determine the GSM fluctuation to observation alphabet mapping. In the context of a HMM we use the distance from the means to discretise the continuous range of signal strength and cell fluctuation. Fluctuation measurements that are close to the activity means indicate a stronger probability of undertaking a particular activity as opposed to those that, in terms of Euclidean distance, are positioned further away. Hence we map these levels of fluctuation to observations contained within the observation alphabet that reflect this likelihood.

In order to learn the levels of fluctuation we collect a series of data points. A data point comprises a cell and signal strength fluctuation measurement. Data points can be collected at random, that is, the cell phone carrier does not

need to declare current activities. In the context of the trial, we can use data from participants who engaged in all activities; walking, driving and remaining still. We do not ask participants to label the data and instead automatically partition the data points into three sets, reflecting the three activities. To partition the data we use K-means with $k=3$, with the assumption that the participant will perform all three activities. If this is not the case then k is adjusted to reflect this.

In Fig. 6a we plot the mean cell and signal strength fluctuations for an urban area on the outskirts of Bristol in the UK. In this figure the amount of fluctuation increases with the speed of the activities, that is, driving produced a greater level than walking, and walking a greater level than remaining stationary. Whilst this relationship is not linear, the positions of these means do lie along an approximately straight line. We have taken approximately 85,000 measurements of cell and signal strength fluctuation from different areas of metropolitan and urban environments and have always found this behaviour to be consistent. That is, the activity means have typically lain along a straight line. On occasion we found the driving mean to rise slightly above the line due to a proportionally greater level of cell fluctuation. Perhaps the most useful aspect of the relationship between the means is that, in terms of Euclidean distance, the driving activity mean will always be greater than the walking activity mean and closer to the walking mean than the still mean and that the still mean will always be closer to and smaller than the walking mean. Hence, given the means produced by K-means, we are able to easily match means to their corresponding activities. This approach would fail if driving produced less fluctuation than walking or if remaining stationary produced a greater fluctuation than walking. We are however happy to take this shortcut because, having conducted extensive experiments in multiple heterogeneous environments we have never found this situation to occur nor do we expect it to. This is due to the nature of GSM handoff strategies and the behaviour of signal strength fluctuation,

i.e. driving creates a greater level of fluctuation than that created by remaining stationary.

Using the mean fluctuation levels learnt for activities in a given environment we are able to define the mappings between GSM cell and signal measurements, and the observation alphabet. We have found the best way to do this is by slicing the 2D measurement space up using the means and variances of the three activities along the two dimensions. An example of this is illustrated in Fig. 6b. Each zone created by slicing the measurement space represents a discrete subsection of the continuous measurement space. Hence each zone represents the membership criteria for an observation. Membership is determined by finding the zone that a current measurement lies in. In Fig. 6b, for clarity, just four observations have been superimposed (v_0, v_1, v_2, v_3). In practice we define 15 observations and subdivide the measurement space into 15 distinct zones.

This approach to learning the optimum settings for a given environment avoids the need to relearn B . Instead we update the mapping between cell and signal strength fluctuation and the observation alphabet. This enables us to provide consistent mappings between GSM measurements and the observation alphabet. This mapping reflects the probability of observations occurring in specific hidden states, matrix B . The alternative approach, relearning B , would require the use of fixed cell and signal strength fluctuation boundaries for mappings to observations. The probability of these observations occurring in specific states for a given environment would then need to be learnt. In order to learn these probabilities would still require the use of K-means to learn the activity variances and means. The primary disadvantage of this approach is that the discretised 2D space that represents the observation alphabet mappings needs to be exhaustive in order to operate in all types of environments. That is, the measurement space needs to be discretised in a manner fine enough to allow operation in environments with low levels of fluctuation as well as those

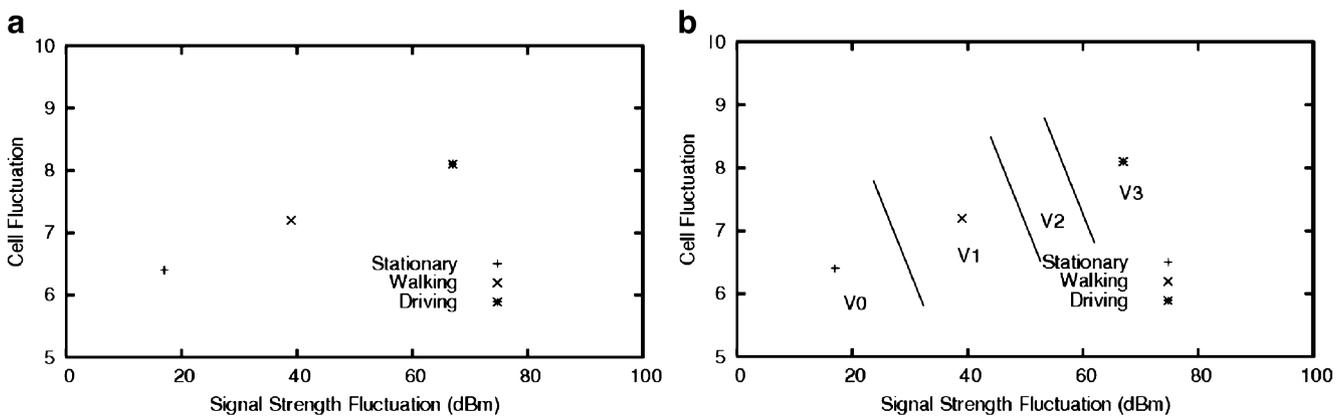


Figure 6 a The activity means and b creating the discrete observation alphabet

with high levels such as metropolitan environments. This requires a much larger observation alphabet. This has a negative impact on the computational overhead of running the Viterbi algorithm to determine the most likely hidden state.

7.3 HMM performance

To assess performance we compared the HMM that was trained using the unsupervised calibration procedure described in the previous section with the Artificial Neural Network (ANN) approach presented in Section 3.1. In order to do this comparison we used the labelled trial calibration data collected before the Shakra user study started.

Both the HMM and the ANN were exposed to approximately 15 min of training data for each activity. As before cell and signal strength fluctuation was measured over 15-s intervals. In the case of the ANN training was conducted by repeatedly presenting data collected during each method of movement. This needs to be carried out on a desktop PC, not directly on the cell phone. The HMM was trained using Baum-Welch to populate A , B and π , and the GSM mappings to the observation alphabet were learnt using the method presented in the previous section. The HMM comprised an observation alphabet of 15 distinct observations. Both the trained ANN and HMM ran on the cell phone in real time.

In order to compare the performance of the ANN and the HMM we presented both algorithms with the same GSM data. We used approximately 2 h test data while undertaking each activity. Data was collected at different times of the day on different days of the week. The results of the ANN and HMM are shown in Table 2.

The ANN did not perform as well as the HMM when sensing if the carrier was stationary. We suspect this is partly due to the nature of the environment. We found signal strength fluctuation to, on occasion, behave in a sporadic manner despite the cell phone carrier remaining stationary. In addition the task knowledge was applied using a simple averaging filter in the ANN whilst a superior 5-step Markov model was used in the HMM.

Whilst walking we found the ANN to perform slightly better than the HMM. We suggest that if the HMM were given more training data a level similar to that of the ANN would be achieved. Inferring that the cell phone carrier is driving is the hardest of the three activities to sense. This is represented in the confusion matrix for the ANN and HMM. The reason that this activity is so hard to sense is due to the nature of driving in metropolitan environments. The self-calibrating HMM performed slightly better than the ANN.

In summary, we have found that a HMM using an unsupervised calibration process to learn the settings for a given environment is able to offer a similar level of performance to that of an ANN that has been manually trained.

Table 2 Confusion matrix

	Stationary	Walking	Driving
ANN (supervised calibration: metro environment)			
Stationary	83%	16%	1%
Walking	5%	87%	8%
Driving	3%	24%	73%
HMM (unsupervised calibration: metro environment)			
Stationary	92%	8%	0%
Walking	12%	80%	8%
Driving	4%	22%	74%

We believe that this simple approach to calibration and modelling task will increase performance and usability even in disparate environments, although further experimentation will be needed to validate this claim, as discussed in following section.

8 Conclusion

The development of Shakra is a first step towards creating a low cost physical activity monitor and health promotion application that is easily accessible to the general public. Shakra's real-time collaborative aspects and its lack of sensors beyond the mobile phone differentiate it from other research and products in this area. The initial reaction to Shakra during its pilot study was very encouraging; however, some issues with accuracy, feedback, privacy and awareness were raised. We have made some first steps towards resolving problems of accuracy, and of training the system to be accurate, with the introduction of unsupervised calibration and task modelling, the remaining issues of feedback, privacy and awareness must be addressed in future implementations.

All of the study participants responded positively towards the system and were tolerant of the momentary lags and jitters in activity classification (as discussed in Section 5). Many of the participants were excited to see their own activity levels, expressing higher motivation and increased awareness. We observed some of the same features that have been seen in more traditional collaboration in exercise to lead to more exercise being done, such as encouragement among 'buddies' and, in some cases, strong competition. The way in which the application was used varied between individuals and groups: it was used variously as a mutual awareness tool, a self-monitoring device and as a game. This highlights the need for a degree of flexibility within the design of a health-promoting system that has a broad user demographic; enabling individuals and groups to use the system in such a way that suits and benefits them.

As with any pilot study, there are limitations to the validity of any resultant claims made. It is not possible, for

example, to claim that over a longer period of time the participants would remain enthused and continue to feel motivated by the system. What we do infer from the pilot study is that Shakra is usable and can initiate such positive responses, and suggest that with further development the system may prove to be an effective health promotion tool.

In addition to improving the accuracy of activity inference, the granularity of activity inference must be increased if we are to more finely categorise the various levels of activity intensity. We believe that this is achievable following the introduction of the new activity detection methods discussed in Section 7. Instead of training the system to recognise walking at any speed, it should be trained at the various intensity cut-off points, e.g. low intensity below 4 mph, moderate intensity above 4 mph. Similarly, work will be done to detect other activities such as cycling, with equivalent distinction between low, moderate, and high intensity cycling. That is not to suggest that a system such as ours that can only distinguish between walking, driving and being stationary is without practical use. Indeed our pilot study showed that this level of granularity was enough to raise awareness and generate discussion about activity levels. Pedometers have been proved to increase the activity levels of users while only monitoring step-count. If any remaining types of activities or contributing factors to intensity are to be acknowledged by the system, or indeed the achievable level of accuracy is deemed inappropriate for long-term use, then we envision the need to utilise additional technology. We intend to explore new sensing and analysis techniques that can run on commodity phones, especially as they evolve to contain such previously exotic hardware as WiFi, GPS and, in phones such as the Nokia 3220, accelerometers.

As the focus of the system is to primarily encourage small changes in behavior, no attention has been made so far to the minimum recommended session length of 10 min. This could be easily introduced by a post-processing 10 min rolling filter, earning users additional accreditation when a 10 min session is completed. Another potential avenue of exploration is that of an adaptive system that evolves alongside a user's activity pattern; the 10 min session accreditation being introduced when the system detects substantial levels of intermittent activity throughout the day.

Once completed the system will be the subject of a clinical trial to determine the extent of any resulting changes in attitude, behaviour and health. We expect to use both qualitative evaluation techniques to assess these changes in objective terms, and qualitative evaluation techniques, to explore the detail of individual and social interaction around the system. In particular, we are interested in how people weave such technology into everyday life [4], and expect users to develop tactics and strategies for use beyond our expectations, appropriating or even 'hacking'

the technology to suit their own goals, desires and contexts. There is clearly great potential in technical explorations using highly accurate specialised assemblies of hardware and software, such as the multimodal sensor board and iMote of [10], but our study illustrated the pragmatic advantages of a lightweight application running on a mobile phone with no such specialised sensors, and no cumbersome attachments, e.g. being strapped to the body. We suggest that a commodity platform will help such a health promoting application be more readily integrated into the lives of the wider population sooner rather than later.

References

1. Adams J, White M (2005) Why don't stage-based activity promotion interventions work? *Health Educ Res* 20(No. 2): 237–243
2. Anderson I, Muller H (2006) Context awareness via GSM signal strength fluctuation. In: *The 4th international conference on pervasive computing, late breaking results*. Oesterreichische Computer Gesellschaft, pp 27–31 May
3. Baum LE, Peterie T, Souled G, Weiss N (1970) A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *Ann Math Stat* 41(1):164–171
4. Bell M, et al. (2006) Interweaving mobile games with everyday life. *Proc ACM CHI*, pp 417–426 (Also available at <http://www.dcs.gla.ac.uk/~matthew/papers/CHIYoshiWithCopyright.pdf>)
5. Consolvo S, Everitt K, Smith I, Landay JA (2006) Design requirements for technologies that encourage physical activity. *Proc ACM CHI*: pp 457–466
6. Department of Health (2004) At least five a week: evidence on the impact of physical activity and its relationship to health. DoH, London
7. Duncan SC, Duncan TE, Strycker LA, Chaumeton NR (2004) A multilevel approach to youth physical activity research. *Exerc Sport Sci Rev* 32(No.3):95–99
8. Ipsos Mori, e-MORI Technology Tracker, December 2002 <http://www.mori.com/technology/techtracker.shtml> [last accessed 29th June 2006]
9. King AC, Friedman R, Marcus B, Castro C, Forsyth L, Napolitano M, Pinto B (2002) Harnessing motivational forces in the promotion of physical activity: the Community Health Advice by Telephone (CHAT) Project. *Health Educ Res* 17(No. 5):627–636
10. Korp P (2006) Health on the internet: implications for health promotion. *Health Educ Res* 21(No. 1):78–86
11. Lester J, Choudhury T, Borriello G, Consolvo S, Landay J, Everitt K, Smith, I. Sensing and modeling activities to support physical fitness. Workshop paper in UbiComp '05 Workshop: Monitoring, Measuring and Motivating Exercise: Ubiquitous Computing to Support Physical Fitness. Tokyo, Japan, 2005
12. Licoppe C, Inada Y (2006) Emergent uses of a Multiplayer Location—Aware Mobile Game: the Interactional Consequences of Mediated Encounters. *Mobilities* 1(No. 1)
13. McLean N, Griffin S, Toney K, Hardeman W (2003) Family involvement in weight control, weight maintenance and weight-loss interventions: a systematic review of randomised trials. *Int J Obes* 27:987–1005
14. Mullen E, Markland D (1997) Variations in self-determination across the stages of change for exercise in adults. *Motiv Emot* 21 (No. 4):349–362

15. Pate RR, Pratt M, Blair SN, Haskell WL, Macera CA, Bouchard C (1995) Physical activity and public health: a recommendation from the Centers for Disease Control and Prevention and the American College of Sports Medicine. *J Am Med Assoc* 273:402–407
16. Rabiner LR (1990) A tutorial on hidden Markov models and selected applications in speech recognition. In *Readings in speech recognition*. Morgan Kaufmann, San Francisco, CA, USA, pp 267–296
17. Siewiorek DP, Smailagic A, Furukawa J, Krause A, Moraveji N, Reiger K, Shaffer J, Wong FL (2003) SenSay: A Context-Aware Mobile Phone. In *Proceedings of 7th International Symposium on Wearable Computers, ISWC*. IEEE Computer Society, pp 248–249
18. Sparling PB, Owen N, Lambert EV, Haskell WL (2000) Promoting physical activity: the new imperative for public health. *Health Educ Res* 15(No. 3):367–376
19. Speck BJ, Harrell JS (2003) Maintaining regular physical activity in women—evidence to date. *J Cardiovasc Nurs* 18(No. 4):282–291
20. Tudor-Locke C (2002) Taking Steps Toward Increased Physical Activity: Using Pedometers to Measure and Motivate. President's Council on Physical Fitness and Sports, June
21. Vandelanotte C, De Bourdeaudhuij I (2003) Acceptability and feasibility of a computer-tailored physical activity intervention using stages of change: project FAITH. *Health Educ Res* 18 (No. 3):304–317
22. World Health Organisation: Move for Health <http://www.who.int/moveforhealth/en/> [last accessed 29th June 2006]
23. Viterbi AJ (1967) Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Trans Inf Theory* 13(2):260–269



Ian Anderson is a Ph.D. student working as part of the Mobile and Wearable Computing Group for the University of Bristol. His research focuses on providing practical context-aware behaviour on mobile devices such as phones. He has applied machine learning techniques to identify aspects of context such as activity and location without using traditional sensors such as accelerometers and GPS devices. Previously he completed a B.Sc. (Hons) in Computing for Real Time Systems at the University of the West of England.



Julie Maitland is a second year Ph.D. student at the Department of Computing Science at the University of Glasgow, where she also

obtained a B.Sc. (Hons) in Software Engineering. Her Ph.D. work explores the underlying dynamics of peer-involvement in health-related behaviour management and investigates how social support can be effectively integrated into the design of health-related behaviour management systems. She comes from a healthcare background; having previously been trained at the University of Manchester's School of Nursing and Midwifery and continues to work as a Staff Nurse in various acute settings throughout Glasgow and Paisley.



Scott Sherwood is a third year Ph.D. student at the Department of Computing Science at the University of Glasgow. Scott has focused upon the presentation of self in every day life and in gaming; both in the physical and digital realms. He is particularly interested in the manner in which information about sign-equipment, such as physical possessions, can be represented in digital media. For example, in games information about scores and a user's inventory is often displayed to other player's – he is interested about how physical information about a user, such as clothing, may be viewed by others in an online environment.



Louise Barkhuus is a Research Assistant at Glasgow University, where she have worked since January 2005. She is affiliated with the Equator group which is a UK wide research project investigating issues in ubiquitous computing. At the moment, she is working on projects involving fitness technologies and activity tracking including focusing on how ubiquitous computing can be used to motivate people to become more active. She also explores the use of entertainment media in the home, in particular TV and video.



Matthew Chalmers is a Reader in Computing Science at Glasgow University, UK. His background is in Computer Science, a B.Sc.

(Hons) at University of Edinburgh, then a Ph.D. at the University of East Anglia in ray tracing and object-oriented toolkits for distributed memory multiprocessors. His current research aims to take account of social and perceptual issues both in the design and theory of computer systems, especially ubiquitous computing systems. He previously worked at Xerox PARC, Xerox Euro-PARC, UBS Ubilab and the University of Hokkaido. He has served on the committees of ACM CHI, Pervasive, Ubicomp and PerCom, and is on the editorial boards of the journals *Pervasive* and *Mobile Computing and Information Visualization*.



Malcolm Hall is a fourth year Ph.D. student at the Department of Computing Science, University of Glasgow advised by Dr. Matthew Chalmers. His research involves examining the ways people share knowledge about new or more efficient software tools and practices, in addition to developing a system architecture that supports this behaviour. Recently, he has worked on building the Domino architecture, a mobile code recommendation system which uses spontaneous ad hoc WiFi networks between mobiles, designed to support incremental adaptation and ongoing evolution of ubiquitous computing systems. He has extensive experience in developing applications for mobile devices including WiFiFoFum, a WiFi scanner used by over a million people and he has consulted on various commercial projects, including Navizon, a hybrid positioning system.



Barry Brown is an innovative interdisciplinary researcher who successfully combines the social and computing sciences. In the last five years, in over 50 peer-reviewed publications (in top forums such as CHI, TOCHI, CSCW and UBICOMP), he has described how computing technology can be better designed using a social science perspective. Dr. Brown has pioneered the serious study of leisure and enjoyment, examining existing leisure practice, new technologies for leisure, and trials of systems in use in real settings. Studying a range of different leisure activities, such as video game playing, tourism and sport, he has applied sociological observations to developing these new technologies. This has pioneered advances such as mixed reality museum visiting, mobile collaborative tourism and augmented-reality video games.



Henk Muller obtained a Doctorate in Computer Science from the University of Amsterdam in the Netherlands and is now a Reader in Computer Science at the University of Bristol. His present research interests include the design and evaluation of computer systems, in particular, light-weight, mobile systems. He heads the Mobile and Wearable Computing group at the University of Bristol and was the group's Principal Investigator for the Equator IRC.