Smartphones for Large-scale Behaviour Change Interventions

Neal Lathia[1], Veljko Pejovic[2], Kiran K. Rachuri[1]
Cecilia Mascolo[1], Mirco Musolesi[2], Peter J. Rentfrow[3]
[1]Computer Laboratory, University of Cambridge, UK
{neal.lathia|kiran.rachuri|cecilia.mascolo}@cl.cam.ac.uk
[2]School of Computer Science, University of Birmingham, UK
{v.pejovic|m.musolesi}@cs.bham.ac.uk
[3]Department of Psychology, University of Cambridge, UK
pjr39@cam.ac.uk

Index Terms

Smartphone Sensing, Distributed Systems, User/Machine Systems, Behaviour Change Interventions, Social and Behavioral Sciences, Social Psychology.

Abstract

The proliferation of smartphones opens a new horizon for behaviour change interventions. Equipped with cutting-edge sensing technology and high-end processors, smartphones can both unobtrusively sense human behaviour and be an ideal platform for delivering feedback and behavioural therapy. In this article, we describe how modern-day smartphones are paving the way for future mobile-based behaviour change applications, and discuss some of our smartphone-based behaviour sensing applications. We also describe a recent and ongoing project, called *UBhave*, that aims to build a general framework for delivering context-specific, smartphone-centric interventions at a large scale. Finally, we outline the open research challenges in this domain for the pervasive computing community.

I. Introduction

As smartphones proliferate throughout society, so too does the opportunity to leverage these devices to study, understand, and positively affect human behaviour. The growing opportunity to research and influence daily lives has two primary catalysts: first, device manufacturers are quickly bringing smartphones onto the market that have ever richer capabilities, in terms of computational power and sensor availability. Typical modern-day smartphones can sense their orientation, acceleration in three dimensions, location, and can record audio. These standard features seamlessly allow researchers to access data streams that reflect the device owner's habits, activities, and routines. In addition to sensors, smartphones today place more digital memory and processing capabilities into individuals' pockets than computers of decades past gave to peoples' desktops. This advance ushers in an era where powerful machine learning algorithms for statistical inferences from sensor data can be designed to run on commodity phones. The second catalyst centres on smartphone owners and the developing culture around smartphone usage, rather than the devices themselves: they have become indispensable to many peoples' daily lives. Their continuous presence and usage is what allows researchers to link the sensor data they can collect back to the owner of the device. Beyond being present – or indeed, within arms reach – for large proportions of their owners' day, smartphones are increasingly being used as the main device to participate in social networks, query the web, and, more broadly, access and produce information.

These rapid technological developments and the widespread adoption of smartphones raises the question of whether smartphones could provide an effective mechanism for tackling ongoing concerns and challenges related to the global population's health and well-being. The list of problems is far reaching: for example, about 285 million people worldwide suffer from diabetes; in the United States, 36% of the adult population is obese; smoking kills nearly six million people every year. Similarly, about 1 in 6 Americans report a history of depression[1]. In many of these cases, lifestyle changes – brought on by means of informing, teaching and supporting those who seek to change – may lead to positive health outcomes. One of the well-studied means that behavioural scientists have developed to induce these changes are Behaviour Change Interventions, or BCIs. These techniques have recently been brought online, and can now be delivered over the Internet. However, they have yet to be fully ported to and integrated with sensor-enhanced smartphones, where they could be directly linked to the behaviours and patterns of daily life that the phone can infer about its owner. Recent work has shown that inferences about both user contexts and physical activities [1], [6] as well as mental states (including emotions [13] and stress [8]) can be approximated using data from smartphone sensors. Both of these relate directly to contexts where BCIs are being studied; for example, weight-reduction interventions are tied to the physical activity of participants, while interventions related to mood disorders are inherently related to monitoring participants' mental states.

In this article, we examine the open questions and challenges of merging smartphone sensing and BCI applications: our goal is to highlight recent work from the mobile sensing domain and describe how it can support the design of smartphone BCIs, as well as review how the structure of BCIs can be transposed onto, and augmented by, mobile sensing applications. In the following section, we begin with an overview of BCIs, how they are now available online, and a number of shortcomings that they cannot directly solve. We then decompose the procedure of delivering behavioural change interventions into a sequence of behaviour monitoring, learning, and providing tailored information, and we show how recent efforts in the mobile sensing domain have begun to tackle different steps of this cycle. Finally, we describe a current project, called *UBhave*, that seeks to bind these pieces together. We close by highlighting a number of ongoing research challenges that may guide multidisciplinary efforts to join smartphones and sensing research with the behavioural sciences.

II. Digital Behaviour Change Interventions

Behaviour Change Interventions (BCIs) centre around advice, support, and relevant information; traditionally, they have been used to improve both physical and mental well-being. Doctors, therapists, teachers, and coaches, in going about their daily activities, deliver BCIs as they guide their patients and students. Behavioural scientists have used the survey data that BCIs gather to study human behaviour. Naturally, BCIs have a limited reach and scale: they are constrained by the time and costs associated with patients meeting their therapists, and remain inaccessible to those in remote areas. These challenges have been tackled by turning to the Internet as a medium for mass BCI delivery, and developing *Digital* Behaviour Change Interventions (DBCIs) that allow researchers and practitioners to reach an audience that extends well beyond their time and budgetary constraints. DBCIs can provide continuous, multi-modal access to information and surveys, and tools like LifeGuide[2] have been developed to allow for DCBIs to be seamlessly designed, built, and deployed to the Web. These interventions fully automate the role and interactions that patients would traditionally have with their therapist: they can respond with tailored advice to users' answers to questions, support goal setting and help users plan and chart their progress, and send personalised emails or SMS reminders. Providing such potential for high levels of interactivity and availability, paired with their low cost, has been shown to have a positive outcome on a variety of behaviours, including tobacco and substance use, diet, sexual behaviour, and stress [15].

Although they broaden the reach and scale of behaviour change interventions, by automating the process of soliciting and delivering tailored information, Web-based DBCIs are characterised by three major constraints: they have stringent requirements for *when* patients interact with them, they depend

[1] http://www.gallup.com
[2] http://www.lifeguideonline.org/

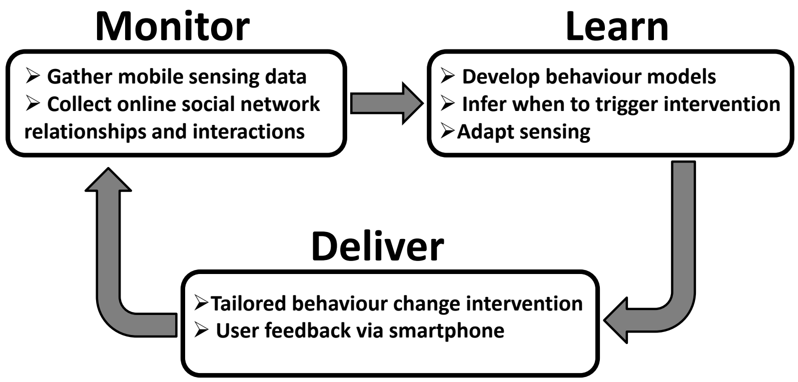

Fig. 1. The three key components of digital behaviour interventions using smartphones: monitor behaviour, learn and infer behavioural patterns, and deliver target behaviour change.

on participants' *self-reporting* to monitor progress, and are typically adopted by self-selecting groups of participants. The former constraint is due to the technical limitations of delivering a DCBI via a Web browser, which typically entails that the intervention can only be delivered to users who are at a computer. While this is certainly less constraining than face-to-face meetings with a therapist, patients may not have access to relevant information in the *moment* that matters the most; for example, accessing dietary advice at meal times, while out at a restaurant. The second constraint similarly stems from the restrictions that browsers impose on collecting data relating to users' behaviours, which must be conducted via surveys and self-reports; this kind of input may be subject to reconstruction bias or inconsistent with the patients' actual behaviours. Finally, participation in DCBIs continues to be primarily for targeted populations: engaging and spreading positive health outcomes throughout society is an ongoing challenge.

Sensor-enabled smartphones are poised to readily solve two of these challenges, and act as a gateway to solve the third. Smartphones have been adopted across the globe, and are regularly used to access local and social information. Recent research surrounding collecting and making inferences about peoples' physical and mental well-being using smartphone sensor data and delivering information via feedback interfaces includes the methodologies and technical solutions to passively monitor users' progress and give them tailored information in the moments they need it the most. Furthermore, DBCIs that leverage inferences from sensor data have the potential to uncover the social and psychological triggers that affect the behaviours that are trying to be changed. For example, knowing that someone who wants to quit smoking has a strong urge to smoke when stressed, while with certain people, or at specific locations, means that the intervention can be delivered at the right time, place, and context. These tasks of the process of behaviour intervention through smartphones are exemplified in Figure 1. In the following sections, we describe a set of systems that support passively monitoring, learning from, and feeding information to smartphone owners about their behaviours: we will then discuss how these can be merged with online social networks and traditional DBCI tools to design future behaviour change applications.

III. Smartphones based Behavioural Monitoring

Smartphones are *ubiquitous*, *unobtrusive*, and *sensor-rich* computing devices. They are carried by billions of users every day; more importantly, their presence is likely to be "forgotten" by their owners, allowing for the passive and effortless collection of data streams that capture user behaviour. Typical modern day smartphones, such as the Samsung Galaxy SIII or the Apple iPhone 5, have a wide range of sensors embedded in them, including an accelerometer, compass, GPS, microphone, and screen proximity sensor, to name a few. Yet the data that is passively available from these devices does not end here: if we also consider that the radios in smartphones are kinds of sensors, then the list increases to include Bluetooth, GSM, Wi-Fi, and NFC. Smartphones are also equipped with powerful processors (such as the

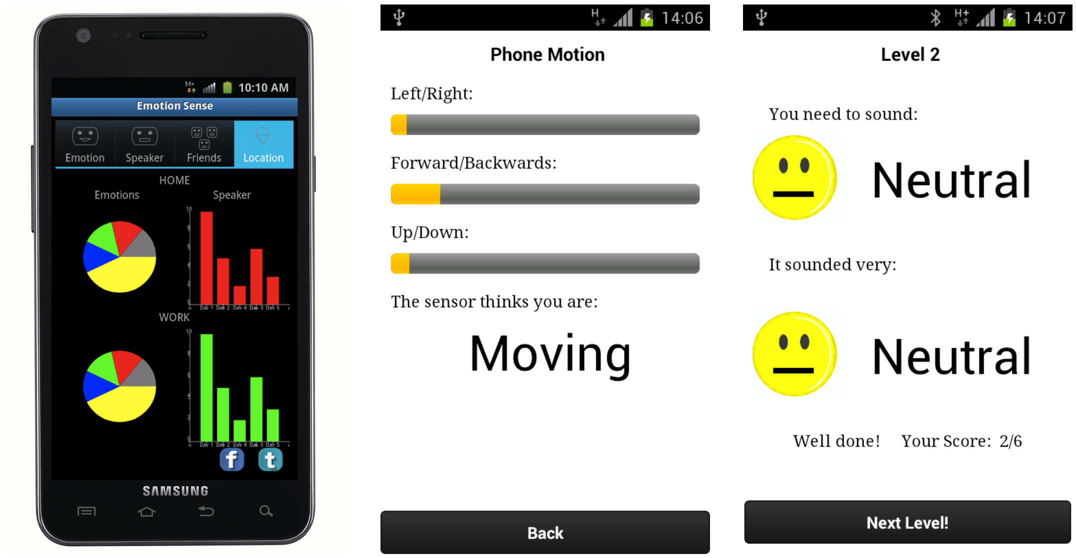

(a) Location Feedback Screen (b) Emotion Feedback Screen (c) Emotion Feedback Screen

Fig. 2. Screenshots of prototype applications that use the emotion inference algorithms. Left: the app used during the trial, which gave feedback about the distribution of inferred emotions. Middle/Right: an app used to demo the inference algorithms, which allows participants to record an audio clip of their voice and receive the inferred emotion (right) and view live feedback about motion detection (middle).

Galaxy SIII Quad-core 1.4 GHz Cortex-A9) that enable them to locally compute intensive classification tasks such as voice processing or image recognition.

The availability of data from these sensors, blended with local computational power and machine learning techniques, allows for smartphones to be used to autonomously infer various activities of the user. For example, physical states such as running, walking, and driving can be inferred from accelerometer data, the microphone sensor can reveal users' conversation status (speaking or not), the Bluetooth radio can be used to detect recurring co-location with other Bluetooth devices (including other phones), and GPS data can track users' locations. More recently, research has shown how similar blends of data, machine learning, and onboard processing availability can be used to make inferences about peoples' mental states, including their emotions and stress levels [8], [13]. How can these systems augment DBCIs? What ongoing challenges are they characterised by?

As described above, current DBCIs rely on users' self-reports to monitor progress and understand users' moods; these self-reports may be subject to reconstruction bias and, moreover, will only be available at the times that users volunteer them. As a potential alternative to mood self-reports and surveys, we designed EmotionSense [13], a passive monitoring smartphone application that can autonomously capture emotive, behavioural and social signals from smartphone owners. There are two key components of this system that allows EmotionSense to automatically (a) recognise who is speaking and (b) what emotions they express by means of classifiers running locally on phones; in the following, we give a brief overview of these subsystems and how they were evaluated.

1) **Speaker Recognition**. The speaker recognition component was implemented using the Hidden Markov Model Toolkit[3]. We use a Gaussian Mixture Model (GMM), a machine learning technique, to capture speech and silence; for speech recognition, we collected approximately 10 minutes of voice data from each of our experiment's participants and generated a 128-component universal background GMM, representing the combined speech data. This model is used to detect ongoing conversations. A complementary GMM *silence model* was also trained using silent audio data; its role is to detect and filter silent audio samples, in order to avoid unnecessary processing in the following emotion recognition component. Finally, we also trained per-user models using audio

[3]HTK, *http://htk.eng.cam.ac.uk/*

data from each of our participants: once a conversation has been detected, these were used to infer who was participating.

2) **Emotion Inference**. The design of the emotion inference component is similar to that of speaker recognition component. We first trained a background GMM representative of all emotions and then generated emotion-specific GMMs. However, instead of collecting training data from the users, we used data from the Emotional Prosody Speech and Transcripts library [7], which is a standard benchmark library in emotion and speech processing research. While this library would allow us to train for up to 14 "narrow" emotions (e.g. cold-anger, hot-anger, and panic), we grouped the classes into 5 "broad" emotions that reflect those used in the social sciences literature: anger, fear, happy, neutral, and sad. Inferring emotional states from microphone audio samples is a multi-stage process: first, a recorded audio file is converted into a vectorial representation of the voice signal over time and then compared with the conversation and silence models. If the audio file contains non-silent data, then it is further processed to compare with each of the user-specific models that we pre-loaded on each of the participants' phones. Finally, the model with highest likelihood of match is assigned as the model of the recorded audio file.

We evaluated the EmotionSense system through several offline micro-benchmark tests and a deployment with 18 participants who recorded their emotions in a daily diary. Our benchmark results showed that the system is able to achieve an accuracy of over 90% for speaker identification and over 70% for broad emotion recognition. The results from the deployment showed that the users exhibit neutral emotion far more than the other emotions, and at a high-level the distribution of emotions detected by our system matched the distribution reported by the participants through self-reports. Further, the results also showed that the users exhibit sad and anger emotions lesser in larger groups than smaller groups. These results agree with research in the social psychology literature [4], [9].

The results from our first EmotionSense trial demonstrate the potential that the combination of passive sensor data collection and machine learning have to provide continuous monitoring of participants' emotional states, while also collecting data that is representative of each person's social interactions and mobility. Broadly speaking, it facilitates the collection of data by social psychology researchers, as the system automatically captures and classifies the user activities. This can be used not only to understand the correlation and the impact of interactions, activities, co-location and location on the emotions and behaviour of individuals, but also paves the way for sensor-enhanced DBCIs. As per the example in the previous section, mobile DBCIs that use EmotionSense would be able to deliver information in meaningful moments, and trigger advice, support, and tailored feedback based on both participants' physical and expressed emotional states.

Finally, systems like EmotionSense are currently primed for data collection, rather than feedback and user interaction. In the following section, we review a recent system that steps into this domain.

IV. Smartphone Feedback for Behavioural Change

While smartphones are powerful gateways to sensor data, they are also an ideal platform for providing feedback and interventions to users. Smartphones are personal devices that users spend a significant amount of their time interacting with. Moreover, smartphones enable interventions to be built using inferences from sensor data streams as *triggers* for information delivery, thereby customising them the times they may be most effective. For example, if a user is more likely to comply with an intervention when she is at home, then the GPS sensor can be used as to detect these moments.

We designed *SociableSense* [12], a smartphones based platform for providing realtime feedback to users, in order to foster social interactions and improve their relations. While providing feedback about social interaction may seem to be a rather simple setting, we note that sociability is a key behaviour that changes in patients who suffer from a variety of disorders, ranging from the autistic spectrum to depression. The system we developed was not tested in a medical setting; however, it demonstrates the potential for smartphone technology to monitor facets of behaviour that are directly linked to BCI

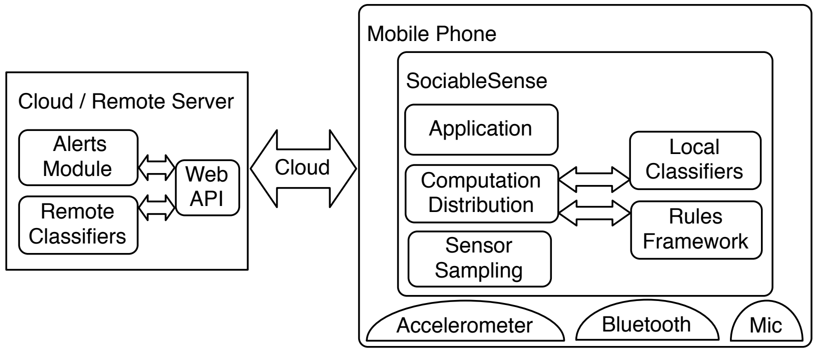

Fig. 3. SociableSense architecture: the data processing and inference tasks are distributed across mobile devices and a cloud-based back-end.

domains. SociableSense, like EmotionSense, captures data from the sensors in off-the-shelf phones. The system then uses this data to model the "sociability" of users, based on their collocation and interaction patterns: it then *closes the loop* by providing real-time feedback and alerts with the aim to make people more sociable. SociableSense relies on the distribution of the computation across mobile devices and a cloud-based back-end as exemplified by the architecture shown in Figure 3.

We define the *sociability* of a user as the strength of the user's connection to his/her social group. The system measures the strength of a user's relations and his/her overall sociability based on the *network constraint* [2]. In a social network, the network constraint for a node quantifies the strength of that node's connectivity to others. For any two persons in a social network, the person with lower *network constraint* value is considered to have higher strength in terms of connectivity and is considered to be more sociable. The SociableSense prototype system measures two relation graphs, based on collocation (the *collocation network constraint*) and interaction (the *interaction network constraint*). We define *collocation* of a pair of users as being in proximity to each other, and *interaction* as speaking, in person, to one another. The system captures collocation patterns through a coarse-grained Bluetooth-based indoor localisation feature. Two sets of Bluetooth devices, representing social locations (such as common rooms, coffee rooms, games/entertainment zones) and work locations (office spaces, meeting rooms, video/audio conferencing rooms) need to be installed in the deployment environment, and by mapping the Bluetooth MAC addresses and locations the system will be able to identify whether a user is in work or social location. This methodology also helps in reducing energy consumption by avoiding Wi-Fi or GPS, that are generally expensive in terms of energy consumption; also, GPS does not generally work in indoor office locations. Interaction patterns, instead, are captured via the microphone sensor and a speaker identification classifier, as described in the previous section.

In order to provide implicit incentives to the users to become more sociable, we implemented a gaming feature that infers the most sociable person in terms of these two constraints, who we refer to as the *mayor* of the group. When a user is in a sociable location, an alert about this is sent to all other participants, so that interested people can join the user and socialise with him/her. The application then displays the strengths of the user's relations and the mayors of the group in order to encourage active participation of the users in the experiment and also to motivate them to socialise.

We evaluated the social feedback component of the SociableSense system with a deployment over two working weeks. In order to understand the effect of the feedback mechanisms and alerts, we conducted the evaluation in two phases: in the first phase, the feedback mechanisms were disabled; in the second phase, they were enabled. Each phase lasted for a week and we measured users' average sociability in terms of the collocation and interaction network constraints in each of these phases. The results showed that the average network constraint with respect to collocation and interaction networks is lower when feedback mechanisms were enabled, i.e., the sociability of the users increased when the feedback mechanisms were displayed. The results also showed that the feedback mechanisms had greater effect in the sociable locations than in the work locations, which might be because of more opportunities to interact. Overall, this deployment has shown that smartphones are a viable platform for providing interventions and feedback to the participants; moreover, phones can also be used to monitor the effect of these interventions on the users. The mechanisms implemented in SociableSense can be used as building blocks for more advanced DBCIs.

V. Closing the Loop: Binding Sensors and Interventions

A new multidisciplinary project, *UBhave*[4], has recently begun to tackle the challenges of bringing mobile sensing into DBCIs. The project is a collaboration between the Universities of Cambridge, Birmingham, Southampton, Oxford, and University College London, in the United Kingdom. The project aims to devise the first holistic platform for large scale DBCI design and delivery, and will focus on and extend the three major components, *monitor*, *learn* and *develop* that we have described above (See Figure 1):

- **Design and Build**. Drawing from the principles that guided the design of the LifeGuide system, the UBhave project aims to build a platform that will make the rapid prototyping, build, and deployment of mobile-based DBCIs, that leverage the power of smartphone sensors, as seamless as possible. Researchers and practitioners whose intent is to build a sensor-enhanced DBCI should not need to know about sensor sampling control and smartphone battery life management: our sensing framework will automate these controls, while allowing those with technical expertise to transparently test their own designs.
- **Monitor and Infer**. The project will harness the EmotionSense and SociableSense frameworks to uncover and expose targeted features of human behaviour that can be extracted from mobile sensors and used to monitor and infer peoples' moods, physical activities, and social relations.
- **Adapt and Learn**. The sensor monitoring components will also be paired with experience sampling questions, daily goals, and surveys. This will enable the system to further learn about how individuals interact with their device in varying contexts, thus allowing for DBCIs to become personalised to the individual who carries the device.
- **Tailor and Deliver**. Bringing smartphone-DBCIs to their full potential will not only mean learning about users via sensors and their feedback: it also entails understanding and detecting when the right moment to deliver tailored information is occurring, and how best to provide feedback to participants.
- **Share and Diffuse**. Finally, UBhave aims to integrate with Online Social Networking sites (OSNs) in order to enhance both the extent it can measure participants' social context, but also to recruit users. Mobile sensing has been integrated with OSNs before [11], but the sensing component was predominantly used to enhance OSN experience. OSNs, however, represent an untapped resource in designing and delivering behaviour change interventions. But OSNs are not just a great way of application diffusion. With such a large number of virtually interconnected users, OSNs are an irreplaceable source of information on psychological characteristics of different groups – essentially another behavioural sensor. In UBhave we will go beyond monitoring of a single user and correlate information gathered from users social network in order to better understand behaviour. Behavioural problems often appear as the result of an impact of the social environment. For example, becoming obese is more likely if one is socialising with obese people [3]. Thus, in UBhave we plan to devise efficient interventions by concentrating on a whole social clique, rather than on an individual. UBhave will deal with data coming from multiple sources: smartphone sensors, OSNs, and event-triggered user surveys. Behavioural modelling under such a diverse set of data sources, and for a large number of users in parallel, is an open problem for statistics and machine learning.

VI. Ongoing Challenges

Starting from these design goals, we now discuss a number of open challenges in mobile sensing for DBCIs, which are also indeed promising research areas for the pervasive computing community.

[4]http://www.ubhave.org

1) **Energy constraints**. Continuous sensing enables real-time monitoring of user's behaviour, yet it quickly depletes mobile device's battery, rendering any DBCI that relies on it unusable. In EmotionSense, energy efficiency is tackled through a dynamic adaptation of sensor duty-cycling that increases the sampling rate when the user's context is changing. In SociableSense we examined the trade-off between sensor sampling frequency and accuracy versus the latency incurred. In order to appropriately monitor and deliver DBCIs, how should sensor sampling adaptation be designed? Will a generic design suffice, or will sensor sampling control become inherently tied to the domain of the DBCI? To what extent can this problem be solved outside of the domain of DBCIs? Or, will they need to be tailored to individual scenarios? For example, does the sensing of activities required for smoking cessation differ so substantially from monitoring emotional well-being that different sensor control techniques will need to be adopted?
2) **Data Processing and Inference Challenges**. The amount of data generated by mobile sensing can quickly surpass the storage and processing capabilities of today's most powerful computers [10]. Ultimately, raw sensing data has to be reduced to features of interest, that cannot only explain human behaviour but also provide actionable recommendations or, at least, be turned into interpretable feedback. Novel statistical tools will have to be developed to extract behavioural features from a large set of heterogeneous data coming from smartphone sensors and online social networks. The intelligent distribution of the computation across multiple heterogeneous devices (mobile and fixed, including a cloud-based back-end) is another open research area.
3) **Generalisability**. People neither behave nor express their behaviour in a uniform way. To what extent do the machine learning techniques that are used, for example, to infer emotional states need to be tailored to individual users? Recently, a number of methods that capture similarity in human behaviour, such as eigenbehaviors [5] and community similarity networks [6], have been proposed. The underlying idea is to identify populations that can be treated as uniform for the sake of behaviour inference. In addition, before delivering DBCIs we need to understand how individual traits and personal attitudes may impact their effectiveness.
4) **Privacy**. Data captured from smartphones sensors raises privacy concerns: for example, microphone recordings used to identify speakers may contain sensitive audio data, and information captured from the GPS sensor may contain locations that users may not want to share. There have been some works on the privacy aspects of smartphone sensing, such as preserving the anonymity of sensor reports without reducing the precision of location data [14]. However, it is possible that systems like the above may record voices of people who have not given their informed consent, which is not permitted in some countries (e.g., the United States). Robust methodologies for smartphone-based speaker identification would begin to overcome this issue; yet, speaker identification is prone to inaccuracies when environmental noise varies and therefore this may not solve the problem entirely. How can systems avoid recording voices of people collocated with the phone user? Moreover, as demonstrated in the SociableSense system, it is more efficient to process audio files in the cloud than locally on the phone. In this scenario, a challenge that needs to be addressed is the implementation of privacy-preserving techniques for remote processing.
5) **Designing for Participation**. Passive sensing of human behaviour through smartphones is possible, as demonstrated by the EmotionSense and SociableSense applications. However, a number of challenges and unanswered research questions remain when it comes to the behaviour intervention aspect. For instance, neither of these, nor any of the other research systems in the area of mobile sensing and DBCIs, solve the problem of tailored, timely intervention and sensing at a large scale. What is the right moment for providing an intervention? What is the effect of various triggers on the user compliance? With smartphone-based DBCIs, users can receive interventions at the most appropriate time and place. Yet, context-specific dissemination of DBCIs remains an unexplored area of research. What role do social and mobile digital technologies play for successful behaviour change interventions? We plan to explore the technological, psychological and social factors influencing uptake, usage and effectiveness of different intervention characteristics and components in promoting

positive behaviour change. We hope to understand when and in which context people do get the most from DBCIs, but also how to engage people in partaking in DBCIs.

6) **Evaluation of sensing-based DBCIs**. Assessing whether a mobile phone application helps in delivering a desired behaviour change intervention has two important questions. First, is the context sensed accurately and with high enough granularity? Sensors in smartphones were not originally designed as a means of capturing behaviour: for example, the microphone sensor is designed for phone calls, rather than speech or emotion recognition. The data that sensors do capture may be inaccurate, not only due to the sensor itself, but also the location of the phone relative to speakers, environmental noise, the cultural background of participants, and their varying emotional expressivity. Second, the evaluation of a DBCI has to answer such complex questions as whether the induced behaviour change is a long-term one, or just a temporary state. In [15], Webb et al. point out the importance of sound psychological theories for successful DBCIs; therefore, we envisage that future systems are likely to require a highly interdisciplinary presence when being designed, built and evaluated. With smartphone-based DBCIs, users can receive interventions at the most appropriate time and place. Yet, context-specific dissemination of DBCIs remains an unexplored area of research. What role do social and mobile digital technologies play in delivering successful behaviour change interventions?

The field of digital behaviour change interventions is rapidly transforming the way we think about wellbeing improvement. Building upon a solid existing body of work on mobile behaviour sensing and Internet-based intervention dissemination, with UBhave, we strive towards comprehensive, large scale DBCIs. We are set to create tools, methods and support systems that will enable a wide inter-disciplinary community to participate in advancing knowledge and skill in the field of digitally supported behaviour change by creating and implementing their own digital behaviour change interventions. Progress in this field of research is currently severely impeded by the fragmented, laborious process of creating individual applications; we will rectify this situation by providing an extensible open source software platform that allows computer science and social science researchers to rapidly develop and easily adapt and share mobile DBCIs.

Acknowledgments

This work was supported through the EPSRC grant "UBhave: ubiquitous and social computing for positive behaviour change" (EP/I032673/1).

References

[1] N. Aharony, W. Pan, C. Ip, I. Khayal, and A. Pentland. The Social fMRI: Measuring, Understanding, and Designing Social Mechanisms in the Real World. In *ACM Ubicomp'11*, 2011.

[2] R. Burt. *Structural Holes: The Social Structure of Competition*. Harvard University Press, 1995.

[3] N. A. Christakis and J. H. Fowler. The Spread of Obesity in a Large Social Network over 32 Years. *New England Journal of Medicine*, 357(4):370–379, 2007.

[4] L. Clark and D. Watson. Mood and the Mundane: Relations Between Daily Life Events and Self-Reported Mood. *Journal of Personality and Social Psychology*, 54.2:296–308, 1988.

[5] N. Eagle and A. Pentland. Eigenbehaviors: Identifying Structure in Routine. *Behavioral Ecology and Sociobiology*, 63:1057–1066, 2009.

[6] N. D. Lane, Y. Xu, H. Lu, S. Hu, T. Choudhury, A. T. Campbell, and F. Zhao. Enabling Large-scale Human Activity Inference on Smartphones using Community Similarity Networks (CSN). In *ACM UbiComp'11*. ACM, 2011.

[7] M. Liberman, K. Davis, M. Grossman, N. Martey, and J. Bell. Emotional Prosody Speech and Transcripts, 2002.

[8] H. Lu, M. Rabbi, T. Choudhury, D. Gatica-Perez, and A. Campbell. StressSense: Detecting Stress in Unconstrained Acoustic Environments using Smartphones. In *UbiComp '12*. ACM, 2012.

[9] M. R. Mehl and J. W. Pennebaker. The Sounds of Social Life: A Psychometric Analysis of Students' Daily Social Environments and Natural Conversations. *Journal of Personality and Social Psychology*, 84(4):857–870, 2003.

[10] G. Miller. The Smartphone Psychology Manifesto. *Perspectives on Psychological Science*, 7:221–237, 2012.

[11] E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng, and A. T. Campbell. Sensing Meets Mobile Social Networks: the Design, Implementation and Evaluation of the CenceMe Application. In *SenSys'08*. ACM, 2008.

[12] K. K. Rachuri, C. Mascolo, M. Musolesi, and P. J. Rentfrow. SociableSense: Exploring the Trade-offs of Adaptive Sampling and Computation Offloading for Social Sensing. In *MobiCom'11*. ACM, 2011.

[13] K. K. Rachuri, M. Musolesi, C. Mascolo, P. J. Rentfrow, C. Longworth, and A. Aucinas. EmotionSense: A Mobile Phones based Adaptive Platform for Experimental Social Psychology Research. In *UbiComp'10*. ACM, 2010.
[14] K. P. Tang, P. Keyani, J. Fogarty, and J. I. Hong. Putting people in their place: an anonymous and privacy-sensitive approach to collecting sensed data in location-based applications. In *CHI '06*, pages 93–102. ACM, 2006.
[15] T. L. Webb, J. Joseph, L. Yardley, and S. Michie. Using the Internet to Promote Health Behavior Change: A Systematic Review and Meta-analysis of the Impact of Theoretical Basis, Use of Behavior Change Techniques, and Mode of Delivery on Efficacy. *Journal of Medical Internet Research*, 12(1), 2010.

Neal Lathia is a Research Associate in the Computer Laboratory, University of Cambridge. His research falls at the intersection of data mining, mobile systems, and personalization/recommender systems. Dr Lathia has a PhD in Computer Science from University College London. More details available at: http://www.cl.cam.ac.uk/∼nkl2.

Veljko Pejovic is a Postdoctoral Research Fellow at the School of Computer Science at the University of Birmingham, United Kingdom. His research is focused on adaptive wireless technologies and their impact on the society. He received a PhD in Computer Science from the University of California Santa Barbara. Contact him at: v.pejovic@cs.bham.ac.uk.

Kiran K. Rachuri is a PhD student in the Computer Laboratory, University of Cambridge. His research interests include smartphone sensing systems, energy efficient sensing, and sensor networks. He received M.S. in Computer Science from the Indian Institute of Technology Madras. Contact him at: kiran.rachuri@cl.cam.ac.uk.

Cecilia Mascolo is a Reader in Mobile Systems in the Computer Laboratory, University of Cambridge. Her interests are in the area of mobility modelling, sensing and social network analysis. Dr Mascolo has a PhD in Computer Science from University of Bologna. More details at www.cl.cam.ac.uk/users/cm542.

Mirco Musolesi is a Senior Lecturer in the School of Computer Science at the University of Birmingham, United Kingdom. His research interests include mobile sensing, large-scale data mining and network science. Dr Musolesi has a Ph.D. in Computer Science from University College London. http://www.cs.bham.ac.uk/∼musolesm.

Peter J. Rentfrow is a Senior Lecturer in the Psychology Department at the University of Cambridge. His research focuses on behavioural manifestations of personality and psychological processes. Dr Rentfrow earned a PhD in Psychology from the University of Texas at Austin. Contact him at: pjr39@cam.ac.uk.