Smartphones for Large-Scale Behavior Change Interventions

Equipped with cutting-edge sensing technology and high-end processors, smartphones can unobtrusively sense human behavior and deliver feedback and behavioral therapy. The authors discuss two applications for behavioral monitoring and change and present UBhave, the first holistic platform for large-scale digital behavior change intervention.

Neal Lathia
University of Cambridge

Veljko Pejovic
University of Birmingham

Kiran K. Rachuri and Cecilia Mascolo
University of Cambridge

Mirco Musolesi
University of Birmingham

Peter J. Rentfrow
University of Cambridge

As smartphones proliferate throughout society, so too does the opportunity to leverage these devices to study, understand, and positively affect human behavior. The growing opportunity to research and influence people's daily lives has two primary catalysts. First, device manufacturers are quickly developing smartphones with rich capabilities in terms of computational power and sensor availability. Second, the developing culture around smartphones focuses on usage rather than the devices themselves, which many users consider indispensable. (For more information, see the "Smartphone Advances Present Opportunities" sidebar.)

Rapid technological developments and the widespread adoption of smartphones raises the question of whether smartphones could provide an effective mechanism for tackling ongoing challenges related to the global population's health and well-being. For example, in the US, 26 percent of the adult population is obese,[1] and smoking kills nearly six million people every year.[2] Similarly, about one in six Americans report a history of depression.[3] In many of these cases, lifestyle changes—brought on by informing, teaching, and supporting those who seek to change—might lead to positive health outcomes.

Behavior-change interventions (BCIs) are a well-studied means that behavioral scientists have developed to induce these changes. 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. Recent work has shown that we can make inferences about user contexts, physical activities,[4,5] and mental states (including emotions[6] and stress[7]) using data from smartphone sensors, which has great potential for BCIs. For example, weight-reduction interventions are tied to the participants' physical activity, while interventions related to mood disorders are inherently related to monitoring participants' mental states.

Here, we examine open questions and challenges related to merging smartphone sensing and BCI applications. We highlight recent work from the mobile sensing domain and describe how it can support the design of smartphone BCIs. In particular, we present two applications for behavioral monitoring and change—EmotionSense and SensibleSense, respectively—and

Published by the IEEE CS ■ 1536-1268/13/$31.00 © 2013 IEEE

Smartphone Advances Present Opportunities

Typical modern-day smartphones can sense their orientation, acceleration in three dimensions, and location, and they can record audio. These standard features let researchers seamlessly 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 in individuals' pockets than computers of decades past placed on peoples' desktops. This advance ushers in an era in which powerful machine-learning algorithms for statistical inferences from sensor data can be designed to run on commodity phones.

Furthermore, smartphones have become indispensable to many peoples' daily lives. Their continuous presence and usage let researchers link the sensor data they collect back to the device's owner. 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.

describe the UBhave project, a holistic platform for large-scale BCI.

Digital Behavior Change Interventions

BCIs center on advice, support, and relevant information. Traditionally, they've been used to improve both physical and mental well-being. Doctors, therapists, teachers, and coaches deliver BCIs to patients and students to guide them as they go about their daily activities. Behavioral scientists have used survey data gathered through BCIs to study human behavior.

Naturally, BCIs are limited in their reach and scale, because they're constrained by the time and costs associated with patients meeting with their therapists. Furthermore, they remain inaccessible to those in remote areas. These challenges have been tackled by using the Internet for mass BCI delivery and developing digital BCIs (DBCIs) that let researchers and practitioners reach an audience that extends well beyond their time and budgetary constraints.

DBCIs can provide continuous, multimodal access to information and surveys, and tools such as LifeGuide (www.lifeguideonline.org) have been developed to help researchers seamlessly design and build DBCIs and deploy them on the Web. These interventions automate interactions that patients would traditionally have with their therapist: they can provide tailored advice, support goal setting, help users plan and chart their progress, and send personalized emails or SMS reminders. Providing such low-cost yet high-level interactivity and availability can positively affect a variety of behaviors, including those related to tobacco and substance use, diet, sexual behavior, and stress.[8]

Although Web-based DBCIs broaden the reach and scale of BCIs by automating the process of soliciting and delivering tailored information, they're characterized by three major constraints. First, they have stringent requirements for when patients can interact with them, owing to the technical limitations of delivering a DCBI via a Web browser, which typically means the intervention can only be delivered to users who are at a computer. Although this is certainly less constraining than face-to-face meetings with a therapist, patients might not have access to relevant information in the moment that matters the most—for example, accessing dietary advice while at a restaurant.

The second constraint is that Web-based DBCIs depend on participants' self-reporting to monitor progress. As with the first constraint, this stems from the restrictions that browsers impose on collecting data relating to users' behaviors, which must be conducted via surveys and self-reports. This kind of input might be subject to reconstruction bias or inconsistent with the patients' actual behaviors.

The final constraint is that they're typically adopted by self-selecting groups of participants. Participation in DCBIs continues to be primarily for targeted populations: 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 about 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 when they need it most.

Furthermore, DBCIs that leverage inferences from sensor data have the potential to uncover the social and psychological triggers that affect the behaviors that patients and therapists seek to change. For example, you can better intervene at the right time and place and in the right context if you know that someone who wants to quit smoking has a strong urge to smoke when stressed, when with certain people, or at specific locations. Figure 1 outlines

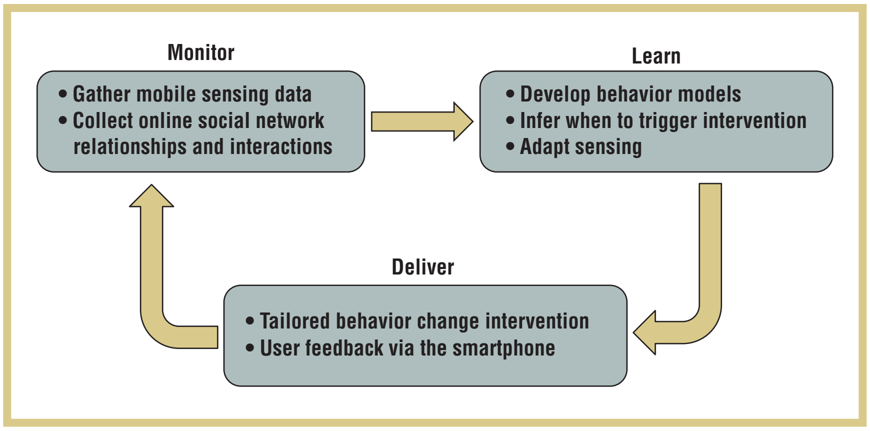

Figure 1. The three key components of digital behavior interventions using smartphones: monitor behavior, learn and infer behavioral patterns, and deliver targeted feedback for behavior change.

the tasks involved with behavior intervention through smartphones.

Behavioral Monitoring

Smartphones are ubiquitous, unobtrusive, and sensor-rich computing devices, carried by billions of users every day. More importantly, owners are likely to "forget" their presence, allowing for the passive and effortless collection of data streams that capture user behavior.

Typical modern-day smartphones, such as the Samsung Galaxy SIII or the Apple iPhone 5, include a wide range of embedded sensors, including an accelerometer, compass, GPS, microphone, and screen proximity sensor. Yet the data that's passively available from these devices doesn't end here. If we also consider the radios in smartphones as sensors, then the list increases to include Bluetooth, the Global System for Mobile Communications (GSM), Wi-Fi, and near field communication (NFC). Smartphones are also equipped with powerful processors (such as the Galaxy SIII Quad-core 1.4 GHz Cortex-A9) that let them 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, lets smartphones autonomously infer various user activities. For example, they can infer physical states, such as running, walking, or driving, from accelerometer data, while the microphone sensor can reveal users' conversation status (speaking or not), the Bluetooth radio can detect recurring colocation 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 make inferences about peoples' mental states, including their emotions and stress levels.[6,7] How can these systems augment DBCIs, and what are the ongoing challenges?

EmotionSense

As we've described, current DBCIs rely on users' self-reports to monitor progress and understand users' moods. Yet such reports can be subject to reconstruction bias and are only available when users volunteer them.

As a potential alternative to mood self-reports and surveys, we designed EmotionSense,[6] a passive monitoring smartphone application that can autonomously capture emotive, behavioral, and social signals from smartphone owners. Two key components let EmotionSense automatically recognize who's speaking and what the speaker is feeling using classifiers running locally on phones.

Speech recognition. We implemented the speaker recognition component using the hidden Markov model Toolkit (http://htk.eng.cam.ac.uk). We use the Gaussian Mixture Model (GMM) machine-learning technique to capture speech and silence. For speech recognition, we collected approximately 10 minutes of voice data from each experiment participant and generated a 128-component universal background GMM, representing the combined speech data. We used this model to detect ongoing conversations.

We also trained a complementary GMM silence model using silent audio data; its role is to detect and filter silent audio samples to avoid unnecessary processing in the emotion recognition component. We also trained per-user models using audio data from each participant: once a conversation has been detected, these models infer who is participating.

Emotion inference. The emotion inference component's design is similar to that of the 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,[9] a standard benchmark library in emotion and speech processing research. Although this library would let us train for up to 14 "narrow" emotions (such as cold-anger, hot-anger, and panic), we instead grouped the classes into five "broad" emotions that reflect those used in the social sciences literature: angry, afraid, happy, neutral, and sad.

Inferring emotional states from microphone audio samples is a multistage process. First, we converted a recorded audio file into a vectorial representation of the voice signal over time, comparing it to the conversation and silence models. If the audio file contains nonsilent data, it's further processed for a comparison with each user-specific model preloaded on the participants' phones. The model with highest likelihood for a match is assigned as the model of the recorded audio file.

EmotionSense Evaluation

We evaluated the EmotionSense system through several offline microbenchmark tests and a deployment with 18 participants, who recorded their emotions in a daily diary. Our benchmark results showed that the system achieved greater than 90 percent accuracy for speaker identification and greater than 70 percent accuracy for broad emotion recognition. The results from the deployment showed that the users exhibited neutral emotion far more than the other emotions, and at a high level, the distribution of detected emotions matched the distribution reported by the participants through self-reports. Furthermore, the results showed that the users exhibited sad and angry emotions less when in larger groups than when in smaller groups. These results agree with research in the social psychology literature.[10,11]

The results from our first EmotionSense trial demonstrate the potential of combining passive sensor data collection and machine learning to provide continuous monitoring of participants' emotional states while collecting data representative of each person's social interactions and mobility. Broadly speaking, approaches such as EmotionSense facilitate the collection of data by social psychology researchers by automatically capturing and classifying user activities. Researchers can use it not only to understand the correlation and the impact of interactions, activities, colocation, and location on the emotions and behavior of individuals but to pave the way for sensor-enhanced DBCIs. Mobile DBCIs that use EmotionSense could deliver information in meaningful moments and trigger advice, support, and tailored feedback based on the participant's physical and expressed emotional states.

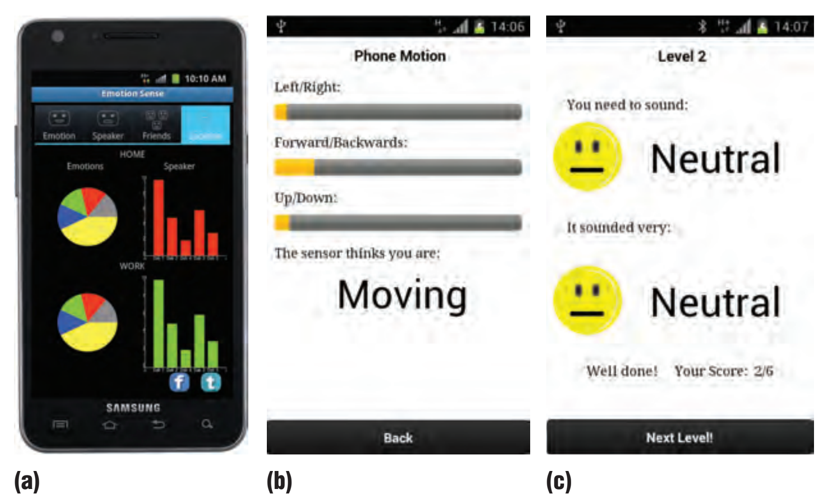

Figure 2. Prototype applications that use emotion inference algorithms. (a) The app used during the trial, which gave feedback about the distribution of inferred emotions. Another app was used to demo the inference algorithms, which lets participants (b) view live feedback about motion detection and (c) record an audio clip of their voice to obtained the inferred emotion.

Behavioral Change

Smartphones are not only powerful gateways to sensor data but also an ideal platform for providing feedback and interventions. Users spend a significant amount of their time interacting with these personal devices. Moreover, smartphones let us build interventions using inferences from sensor data streams as triggers for information delivery, so the interventions are customized for optimal effect. For example, if a user is more likely to comply with an intervention when at home, the GPS sensor detect these moments.

SociableSense

We designed SociableSense,[12] a smartphone-based platform for providing real-time feedback to users to foster and improve social interactions. Although providing feedback about social interaction might seem simple, sociability changes in patients who suffer from a variety of disorders, ranging from the autistic spectrum to depression. The system we developed wasn't tested in a medical setting, but it demonstrates the potential for smartphone technology to monitor facets of behavior directly linked to BCI domains.

SociableSense, like EmotionSense, captures data from the sensors in off-the-shelf phones. It then uses this data to model the users' "sociability" based on their collocation and interaction patterns. The system then closes the loop by providing real-time feedback and alerts to make people more sociable. SociableSense relies on the distribution of the computation across mobile devices and a cloud-based back end (see Figure 3).

We define a user's sociability as the strength of the user's connection to his or her social group. The system measures the strength of a user's relations and his or her overall sociability based on the network constraint.[13] In a social network, the network constraint for a node quantifies the strength of that

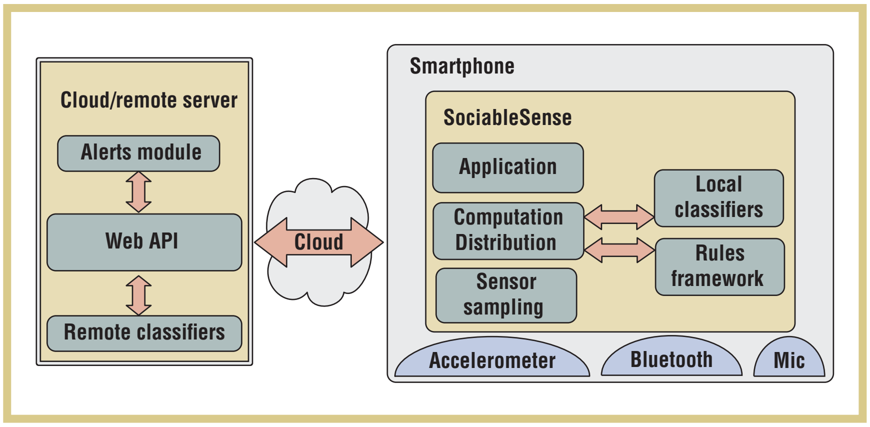

Figure 3. SociableSense architecture: the data processing and inference tasks are distributed across mobile devices and a cloud-based back end.

node's connectivity to others. For any two people in a social network, the person with the lower network-constraint value has the higher connectivity strength and is thus considered more sociable.

The SociableSense prototype 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 we define interaction as speaking, in person, to one another. The system captures collocation patterns through a coarse-grained Bluetooth-based indoor localization feature. Two sets of Bluetooth devices, representing social locations (such as common rooms, coffee rooms, and games or entertainment zones) and work locations (office spaces, meeting rooms, and video or audio conferencing rooms) must be installed in the deployment environment. Then, by mapping the Bluetooth media access control (MAC) addresses and locations, the system can identify whether a user is at work or in a social location.

This methodology also helps reduce energy consumption by avoiding Wi-Fi or GPS, which are generally expensive in terms of energy consumption. Also, GPS doesn't generally work in indoor office locations. Interaction patterns, instead, are captured via the microphone sensor and a speaker identification classifier, as described earlier.

To provide implicit incentives to the users to become more sociable, we implemented a gaming feature that infers the most sociable person, who is then referred to as the group's mayor. When a user is in a sociable location, an alert is sent to all other participants so interested people can join the user. The application then displays the strengths of the user's relations and the mayors of the group to encourage active user participation in the experiment and to motivate users to socialize.

SociableSense Evaluation

We evaluated the social feedback component of the SociableSense system with a deployment over two weeks. To understand the effect of the feedback mechanisms and alerts, we conducted the evaluation in two phases. During the first phase, we disabled the feedback mechanisms; we enabled them for the second phase. Each phase lasted one week, and we measured users' average sociability in terms of the collocation and interaction network constraints during each phase.

The results showed that the average network constraint with respect to collocation and interaction networks is lower when feedback mechanisms were enabled. In other words, user sociability increased when the feedback mechanisms were displayed. The results also showed that the feedback mechanisms had a greater effect in social than in work settings, which might be because they present 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 users. The mechanisms implemented in SociableSense can be used as building blocks for more advanced DBCIs.

Binding Sensors and Interventions

A new multidisciplinary project, UBhave (www.ubhave.org), aims 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 UK. The project aims to devise the first holistic platform for large-scale DBCI design and delivery, focusing on and extending the three major components described in Figure 1: monitoring, learning, and developing.

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, building, and deployment of mobile-based DBCIs that leverage the power of smartphone sensors as seamless as possible. Researchers and practitioners who want to build a sensor-enhanced DBCI shouldn't 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 behavior 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 help the system further learn about how individuals interact with their devices in varying contexts, thus making it possible to personalize DBCIs.

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 and how to deliver the tailored information.

Share and Diffuse

Finally, we hope to integrate UBhave with online social networking (OSN) sites to extend its ability to measure participants' social context and to recruit users. Mobile sensing has been integrated with OSN,[14] but the sensing component was predominantly used to enhance the OSN experience. OSN sites, however, represent an untapped resource in designing and delivering BCIs. They offer an irreplaceable source of information on the psychological characteristics of different groups—essentially acting as another behavioral sensor.

UBhave will go beyond monitoring of a single user, correlating information gathered from a user's social network to better understand behavior. Behavioral problems often appear as the result of an impact of the social environment. For example, you're more likely to become obese if you're socializing with obese people.[15] Thus, with UBhave, we plan to devise efficient interventions by concentrating on a whole social clique rather than just the individual.

Furthermore, UBhave will deal with data from multiple sources: smartphone sensors, OSN sites, and event-triggered user surveys. Behavioral modeling 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.

Ongoing Challenges

These design goals present several open challenges in mobile sensing for DBCIs, which are also promising research areas for the pervasive computing community.

Energy Constraints

Continuous sensing enables real-time monitoring of a user's behavior, yet it quickly depletes the mobile device's battery, rendering any DBCI that relies on it unusable. In EmotionSense, we tackled energy efficiency by dynamically adapting sensor duty-cycling to increase 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.

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 DBCI domain? To what extent can we solve this problem outside of the DBCI domain? Or, will sampling control methods 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 we'll need to adopt different sensor control techniques?

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.[16] Ultimately, raw sensing data must be reduced to features of interest that not only explain human behavior but also provide actionable recommendations (or at least data that can be turned into interpretable feedback).

Novel statistical tools must be developed to extract behavioral features from a large set of heterogeneous data coming from smartphone sensors and online social networks. Another open research area is the intelligent distribution of the computation across multiple heterogeneous devices (mobile and fixed, including a cloud-based back end).

Generalizability

People neither behave nor express their behavior in a uniform way. To what extent do the machine-learning techniques used, for example, to infer emotional states need to be tailored to individual users?

Recently, researchers have proposed methods that capture similarity in human behavior, such as eigenbehaviors[14] and community similarity networks.[5] The underlying idea is to identify populations that can be treated uniformly for the sake of behavior inference. In addition, before delivering DBCIs, we need to understand how individual traits and personal attitudes can impact their effectiveness.

Privacy

Data captured from smartphone sensors raises privacy concerns. For example, microphone recordings used to identify speakers could contain sensitive audio data, and information captured from the GPS sensor might contain locations that users don't 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.[18] However, systems such as those discussed here could record voices of people who have not given their informed consent, which is illegal in some countries (such as the US).

Robust methodologies for smartphone-based speaker identification would begin to overcome this issue, yet speaker identification is prone to inaccuracies when environmental noise varies, so this might not solve the problem. How can systems avoid recording voices of people colocated with the phone user? Moreover, as demonstrated in the SociableSense system, it's more efficient to process audio files in the cloud than locally on the phone, but this requires implementing privacy-preserving techniques for remote processing.

Designing for Participation

Passive sensing of human behavior through smartphones is possible, as demonstrated by the EmotionSense and SociableSense applications. However, many challenges and unanswered research questions remain when it comes to behavior intervention. For example, research systems in the area of mobile sensing and DBCIs have yet to solve the problem of tailored, timely intervention and sensing on a large scale. What is the right moment for providing an intervention? How do various triggers affect 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 research area. What role do social and mobile digital technologies play for successful BCIs? We plan to explore the technological, psychological, and social factors influencing uptake, usage, and effectiveness of different intervention characteristics and components in promoting positive behavior change. We hope to understand when and in which context people get the most from DBCIs and how to encourage people to use DBCIs.

Evaluating Sensor-Based DBCIs

Assessing whether a mobile phone application helps deliver a desired BCI involves two important questions: First, is the context sensed accurately and with high enough granularity? Sensors in smartphones weren't originally designed to capture behavior. For example, the microphone sensor is designed for phone calls—not for speech or emotion recognition. The data that sensors capture might be inaccurate, not only due to the sensor itself but also to 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 must answer such complex questions as whether the induced behavior change is long term or temporary. Thomas Webb and his colleagues point out the importance of sound psychological theories for successful DBCIs.[8] Therefore, we envision future systems that 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 research area. What role do social and mobile digital technologies play in delivering successful BCIs?

The field of DBCIs is rapidly transforming how we think about well-being improvement. Building on a solid body of existing work on mobile behavior sensing and Internet-based intervention dissemination, UBhave strives to offer comprehensive, large-scale DBCIs. We are set to create tools, methods, and support systems that will let a wide interdisciplinary community participate in advancing knowledge and skills in the field of digitally supported behavior change by creating and implementing DBCIs. Progress in this field is currently impeded by the fragmented, laborious process of creating individual applications. We hope to rectify this situation by providing an extensible open source software platform that lets computer science and social science researchers rapidly develop and easily adapt and share mobile DBCIs.

ACKNOWLEDGMENTS

This work was supported through the Engineering & Physical Sciences Research Council (EPSRC) grant, "UBhave: ubiquitous and social computing for positive behavior change" (EP/I032673/1) and the Gates Cambridge Trust.

REFERENCES

1. A. Brown, "In U.S., Obesity Rate Stable in 2012," Gallup, 23 Jan. 2013; www.gallup.com/poll/160061/obesity-rate-stable-2012.aspx.
2. "Media Center: Tobacco," World Health Organization, May 2013; www.who.int/mediacentre/factsheets/fs339/en.
3. B.W. Pelham, "About One in Six Americans Report History of Depression," Gallup, 22 Oct. 2009; www.gallup.com/poll/123821/one-six-americans-report-history-depression.aspx.
4. N. Aharony et al., "The Social fMRI: Measuring, Understanding, and Designing Social Mechanisms in the Real World," *Proc. 13th Int'l Conf. Ubiquitous Computing* (UbiComp 11), ACM, 2011, pp. 445–454.
5. N.D. Lane et al., "Enabling Large-Scale Human Activity Inference on Smartphones Using Community Similarity Networks (CSN)," *Proc. 13th Int'l Conf. Ubiquitous Computing* (UbiComp 11), ACM, 2011, pp. 355–364.
6. K.K. Rachuri et al., "EmotionSense: A Mobile Phones Based Adaptive Platform for Experimental Social Psychology Research," *Proc. 12th ACM Int'l Conf. Ubiquitous Computing* (UbiComp 10), ACM, 2010, pp. 281–290.
7. H. Lu et al., "StressSense: Detecting Stress in Unconstrained Acoustic Environments Using Smartphones," *Proc. 14th Int'l Conf. Ubiquitous Computing* (UbiComp 12), ACM, 2012, pp. 351–360.
8. T.L. Webb et al., "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," *J. Medical Internet Res.*, vol. 12, no. 1, 2010; doi: 10.2196/jmir.1376.
9. M. Liberman et al., *Emotional Prosody Speech and Transcripts*, Linguistic Data Consortium, 2002.

10. L. Clark and D. Watson, "Mood and the Mundane: Relations between Daily Life Events and Self-Reported Mood," *J. Personality and Social Psychology*, vol. 54, no. 2, 1988, pp. 296–308.

11. M.R. Mehl and J.W. Pennebaker, "The Sounds of Social Life: A Psychometric Analysis of Students' Daily Social Environments and Natural Conversations," *J. Personality and Social Psychology*, vol. 84, no. 4, 2003, pp. 857–870.

12. K.K. Rachuri et al., "SociableSense: Exploring the Trade-offs of Adaptive Sampling and Computation Offloading for Social Sensing," *Proc. 17th Annual Int'l Conf. Mobile Computing and Networking* (MobiCom 11), ACM, 2011, pp. 73–84.

13. R. Burt, *Structural Holes: The Social Structure of Competition*, Harvard Univ. Press, 1995.

14. E. Miluzzo et al., "Sensing Meets Mobile Social Networks: The Design, Implementation and Evaluation of the CenceMe Application," *Proc. 6th ACM Conf. Embedded Network Sensor Systems* (SenSys 08), ACM, 2008, pp. 337–350.

15. N.A. Christakis and J.H. Fowler, "The Spread of Obesity in a Large Social Network over 32 Years," *New England J. Medicine*, vol. 357, no. 4, 2007, pp. 370–379.

16. G. Miller, "The Smartphone Psychology Manifesto," *Perspectives on Psychological Science*, vol. 7, no. 3, 2012, pp. 221–237.

17. N. Eagle and A. Pentland, "Eigenbehaviors: Identifying Structure in Routine," *Behavioral Ecology and Sociobiology*, vol. 63, no. 7, 2009, pp. 1057–1066.

18. K.P. Tang et al., "Putting People in their Place: An Anonymous and Privacy-Sensitive Approach to Collecting Sensed Data in Location-Based Applications," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 06), ACM, 2006, pp. 93–102.

the AUTHORS

Neal Lathia is a research associate in the Computer Laboratory at the University of Cambridge. His research falls at the intersection of data mining, mobile systems, and personalization/recommender systems. Lathia has a PhD in computer science from the University College London. Contact him at neal.lathia@cl.cam.ac.uk.

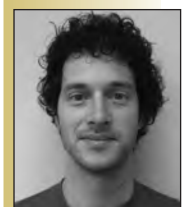

Veljko Pejovic is a postdoctoral research fellow at the School of Computer Science at the University of Birmingham, UK. His research focuses on adaptive wireless technologies and their impact on society. Pejovic received a PhD in computer science from the University of California, Santa Barbara. Contact him at v.pejovic@cs.bham.ac.uk.

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

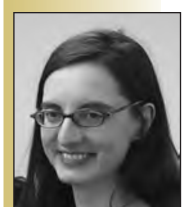

Cecilia Mascolo is a reader in mobile systems in the Computer Laboratory at the University of Cambridge. Her interests are in the area of mobility modeling, sensing, and social network analysis. Mascolo has a PhD in computer science from the University of Bologna. Contact her at cecilia.mascolo@cl.cam.ac.uk.

Mirco Musolesi is a senior lecturer in the School of Computer Science at the University of Birmingham, UK. His research interests include mobile sensing, large-scale data mining, and network science. Musolesi has a PhD in computer science from the University College London. Contact him at m.musolesi@cs.bham.ac.uk.

Peter J. Rentfrow is a senior lecturer in the Psychology Department at the University of Cambridge. His research focuses on behavioral manifestations of personality and psychological processes. Rentfrow earned a PhD in psychology from the University of Texas at Austin. Contact him at pjr39@cam.ac.uk.