
Anticipatory Mobile Computing for Behaviour Change Interventions

Veljko Pejovic

School of Computer Science
University of Birmingham, UK
v.pejovic@cs.bham.ac.uk

Mirco Musolesi

School of Computer Science
University of Birmingham, UK
m.musolesi@cs.bham.ac.uk

Abstract

Behavioural change interventions represent a powerful means for tackling a number of health and well-being issues, from obesity to stress and addiction. In the current medical practice, the change is induced through tailored coaching, support and information delivery. However, with the advent of smartphones, innovative ways of delivering interventions are emerging. Indeed, mobile phones, equipped with an array of sensors, and carried by their users at all times, enable therapists to both learn about the user behaviour, and impact the behaviour through the delivery of more relevant and personalised information. In this work we propose harnessing pervasive computing to not only learn from users' past behaviour, but also predict future actions and emotional states, deliver interventions proactively, evaluate their impact at run-time, and over time learn a personal intervention-effect model of a participant.

Author Keywords

Behaviour Change Interventions, Anticipatory computing, Mobile Computing, Mobile Sensing.

ACM Classification Keywords

H.1.2 [User/Machine Systems]; J.4 [Social and Behavioral Sciences].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
UbiComp '14, September 13 - 17 2014, Seattle, WA, USA '14.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3047-3/14/09\$15.00.
<http://dx.doi.org/10.1145/2638728.2641284>

Introduction

Mobile devices, such as smartphones, represent more than occasionally used tools, and nowadays coexist with their users throughout the day. Therefore, smartphones provide always-on connectivity and the most direct means of information delivery. In addition, these devices are equipped with sensors that can provide momentarily data about a multitude of contextual aspects of the user and the situation, including user's location, movement, social environment and emotional state. The resulting synergy of the user and the smartphone enables not only monitoring and understanding, but also influencing human behaviour. Smartphones' context-awareness and interactivity lead to innovative applications in areas spanning from commerce to public safety. For example, location-relevant ads can be shown to a roaming user looking for a restaurant, alternative routes can be suggested to a driver stuck in traffic, while a careless pedestrian can be alarmed when his phone notices an approaching vehicle [26].

Recently, the smartphone's potential to impact human behaviour has been proposed as a cornerstone of digital behaviour change interventions (dBCIs) [11]. Obesity, depression, alcohol and tobacco abuse are a few of the health and well-being issues that can be successfully tackled through a change in patients' behaviour. DBCIs aim to support those who seek the change, by providing timely and relevant information about the actions that the patient should take. Unlike the traditional therapist-patient interaction, smartphone-based interventions scale to a potentially very large number of users, and can be delivered in accordance to user's momentarily behaviour and state.

While dBCIs are yet to become widespread, we can already argue that the concept of dBCI does not fully

utilise the affordances of mobile computing. In this paper we present our stance that dBCIs should base their awareness of the user not only on the currently inferred context, sensed via smartphone sensors, but also on the future context, predictions of which should be based on the context change models learnt through user and context monitoring. Moreover, we propose to go beyond prediction, and include anticipation – the ability to foresee the effect of current and future actions – when it comes to behaviour change interventions. This concept termed *anticipatory mobile dBCIs* delivers proactive health and well-being interventions, an example of which would be a mobile phone application that can anticipate sedentary behaviour of a person at risk of obesity, and suggests an action, e.g., a joint jogging exercise with a friend.

In the rest of the paper we firstly describe how dBCI can be tailored to, and delivered via mobile phones, with a specific attention paid to smartphone sensing and context awareness. We present the fundamental ideas at the basis of anticipatory computing, and argue that mobile phones represent a suitable platform for its implementation. We then propose anticipatory mobile dBCI, and discuss its key concepts: *context anticipation* and *action-effect modelling*. Many challenges have to be addressed in order to design and implement this paradigm. They include issues in diverse areas of computer science including mobile sensing, behaviour learning and prediction, and human-computer interaction (HCI). Indeed, some of these challenges are fundamental for mobile computing in general (e.g., resource consumption) while some, such as how to foresee and impact futures of mobile users, are more exclusive to anticipatory mobile dBCIs. In this paper we particularly pay attention to the latter obstacles towards the realisation of anticipatory mobile dBCIs. Finally, we conclude with an overview of related and our future work in this domain.

Mobile Behaviour Change Interventions

Smoking and alcohol addiction, stress and obesity are just some of the leading causes of reduced life expectancy that can be successfully tackled with a change in patients' behaviour. Behavioural change interventions (BCIs) are a psychological method that aims to elicit a positive behaviour change. These interventions commonly include collecting relevant information about the participant, setting goals and plans for the participant, monitoring behaviour, and providing feedback in a face-to-face session. BCIs therefore require substantial resources, and the Web has been recognised as distribution tool that can help alleviate some of the requirements. Digital behaviour change interventions (dBCIs) can reach to a participant without a need for a face-to-face contact, can be delivered at any time of day, and can scale to a large number of users. It is worth noting that interventions might not be limited to wellbeing. For example, they might be designed for example to promote "green" living and a more cohesive society.

With more than 5 billion subscribers that carry their devices with them at all times, mobile phones represent the most direct point of contact with almost any individual in the world. Recently, researchers and practitioners have started utilising mobile devices for dBCIs, first via SMS messages [5], and then increasingly via smartphone applications [7]. The affordances of the smartphone, including powerful processing capabilities, a richer interaction interface and an array of sensors – location, proximity, light, acceleration sensors, to name a few – have a potential to not only augment, but completely transform the way behaviour changes are induced.

The key property of smartphone-based dBCIs is that the user monitoring part of the intervention can be automated

through sensing. User movement, physical activity, even stress and emotional state can be inferred from the sensors. This in turn can serve as a basis for reactive timely interventions that will act in a context-aware manner, where the context is provided by smartphone sensing. In BeWell [8], for example, core aspects of physical, social, and mental wellbeing – sleep, physical activity, and social interactions – are monitored via the sensors embedded in the smartphone. The device then provides feedback to the user by showing an aquatic ecosystem where the number and the activity of animals depend on user's wellbeing.

Context inference is a well established research area of mobile computing. Its core lies in machine learning, as raw sensor data from mobile phones are seldom of direct interest. The data is usually processed through sensing, feature extraction and modelling phases, before high-level inferences are made. Increasingly, powerful machine learning algorithms and fine-grained sensor data are being used as a basis for the prediction of future events. Such predictions have been predominantly explored in the mobility domain, where substantial research work has been done on predicting user trajectories, geographical coordinates, and visits to significant places [2].

Initial mobile sensing work did not exploit the fact that phones closely integrate with the user and represent the most direct way of impacting behaviour. In a transition that Campbell and Choudhury call from smart to cognitive phones [4], we can expect smartphones to become proactive personal assistants that anticipate and react upon patterns learnt from our daily life. Such smartphones have the potential to co-exist with their users and narrow the mismatch between the technology and humanity. They will allow for user's movement and

action anticipation; they will know when to attract user's attention and when to stay silent; and they will proactively manage energy and wireless resources. Finally, predictive smartphones will close the loop providing valuable feedback and information that will shape user's future behaviour.

Anticipatory Mobile Computing

An anticipatory system is defined by Rosen as: "A system containing a predictive model of itself and/or its environment, which allows it to change state at an instant in accord with the model's predictions pertaining to a later instant" [19]. Anticipation is deeply integrated with intelligence, and biological systems often base decisions for their actions on predictions. These predictions are based on models that evolve over time. We argue that smartphones fulfil the necessary prerequisites for anticipatory computing: with the help of built-in sensors phones are context-aware, they are capable of fast information processing necessary for predictions, and through interaction with the user they can impact the flow of future.

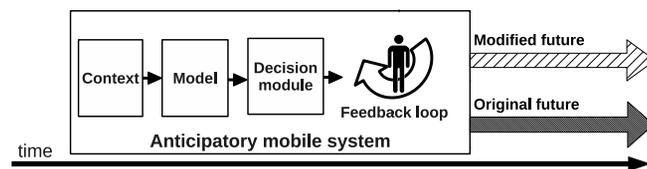


Figure 1: Anticipatory mobile systems predict context evolution and the impact their actions can have on the predicted context. The feedback loop consisting of a mobile and a human enables the system to affect the future.

An anticipatory system relies on the prediction of its environment, thus embraces some of the ideas expressed at the end of the last section. However, *Participation involves more than merely prediction – it requires the understanding*

of how actions within the system might affect the future. In Figure 1 we sketch an anticipatory mobile system that senses the context and builds a model of the environment evolution, which gives it the original predicted future. The system then evaluates the possible outcome of its actions on the future. An action that leads to the preferred modified future is realised through the feedback loop that involves interaction of the system with the user.

Although anticipatory systems have been attempted before, predominantly in robotics [23], the smartphone represents a natural platform for such a system. Not only can smartphones sense and infer the context, build models of the context evolution, but with smartphones the interactive component of anticipatory computing is a reality. Users rely on smartphones for a range of tasks including communication, navigation, and information search. Besides prediction, anticipation requires the ability to filter out irrelevant aspects of the context, and intuitively discard actions that are likely to lead to unfavourable outcome. Humans are well versed in holistic reasoning and picking out relevant information. With smartphones we can establish a bond between human and artificial anticipation, harness phone's sensing, computing and communication abilities to augment human reasoning, but also bootstrap machine learning models with human knowledge.

Anticipatory Mobile dBCIs

We argue that behaviour change interventions can be made timelier and more efficient if delivered through an anticipatory mobile computing system, and term such interventions *anticipatory mobile dBCIs*. In traditional BCI's self-reported user's experiences guide the intervention tailored by the therapist. In web-based and mobile dBCIs, interventions are tailored on-the-fly, responding to the immediate user's context. In anticipatory dBCIs,

however, interventions are selected based on the predicted context. This proactive reasoning pushes anticipatory dBCIs further into the area of preventive healthcare.

We now give an example of an anticipatory dBCI mobile application that proactively tackles depression. Decreased movement, the lack of socialising, and irregular sleep patterns are characteristic symptoms of depression. Depressed people often exhibit sedentary lifestyle [20], and phone's accelerometer, GPS and Bluetooth sensors can be used to detect such behaviour. But even beyond inference, our app could embrace an existing solution to predict the level of user mobility; some of the mobility prediction approaches are surveyed in [2]. Smartphones can also gauge socialisation via Bluetooth sensing [17]. In our app, we could monitor Bluetooth contacts, build a "socialisation" model and then spot periods when the user is predicted to be both alone and not mobile. In addition, we could harness existing sleep pattern inference models (such that used in BeWell), and extend it to predict poor sleeping habits. Once depression is predicted, our app selects and uses a behavioural intervention tool, for example shows a link to a pair of discounted theatre tickets, incentivising the participant to go out and socialise. As the intervention progresses, our app learns about the impact of different tools on user's behaviour, e.g., a change in mobility after a certain message shown, and adjusts the therapy to the specific user.

Implementation Issues and Open Research Challenges

Rooted in mobile sensing, anticipatory mobile dBCI face challenges such as energy inefficiency of continuous sensing and reliable classification of sensed phenomena. On the behavioural intervention side, further experiments are needed in order to evaluate the benefits, if any, of mobile dBCIs. A thorough discussion of challenges in

mobile sensing, anticipatory mobile computing and mobile interventions is available elsewhere [9, 15, 7], therefore, in the rest of the section we concentrate on aspects that are unique to anticipatory mobile dBCIs.

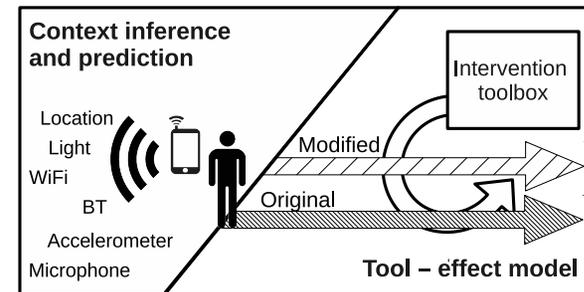


Figure 2: Machine learning is used to infer and predict user context from phone's sensors, but also to create a model of how actions from the behaviour intervention toolbox impact the health and well-being state of the user.

Learning with the User. Machine learning techniques play a major role in the design of anticipatory mobile systems. In Figure 2 we sketch two major aspects in which machine learning supports anticipatory mobile dBCIs. On one side, machine learning is used to infer and predict context in which a user is, as well as the user's internal context, from phone usage and mobile sensor data. Rapid advancements of mobile computing, including an ever increasing sensing capabilities of smartphones promoted a flurry of research that deals with context inference and prediction, ranging from activity recognition to the forecast of user's next location [22], and to the inference of human emotions [18], stress levels [12] and interruptibility [16]. Being able to accurately predict the evolution of context and behaviour is crucial for anticipatory dBCIs to function properly. On the other side

of Figure 2 we have a machine learning model of the impact of interventions on predicted user behaviour. This model is labelled *tool-effect model*, and aims to uncover the relationship between available behaviour change intervention tools, such as suggestions shown to a user via smartphone and the induced change in behaviour, which can be detected via smartphone sensors. This concept is unique to anticipatory mobile dBCIs, and is yet to be explored in the area of mobile computing, and in this paper we discuss means of realising tool-effect modelling.

Reinforcement Learning. The key issue in tool-effect modelling is mapping the space of behaviour change induced by each of the BCI tools, so that tools leading to improved behavioural outcome are identified. Reinforcement learning where an agent uses BCI tools in the intervention environment (which for example can be represented through a Markov decision process) is a natural way to model the problem [24]. In every step, a decision to use a certain tool comes with a change in behaviour and elicit a reward that reflects how positive the change is. To bootstrap the learning process, the initial tool selection logic could rely on the psychiatric practice – an expert therapist would propose tools according to the target aspect of behaviour, the context in which a tool can be fired and even the traits of a specific patient.

Measuring Behaviour. To model the tool-effect loop we need to evaluate the effect of an intervention action, i.e., compare the *original predicted behaviour* and the *modified predicted behaviour*. The metric used for comparison depends on the purpose of the intervention. For example, if dealing with obesity, one possible behaviour metric can be the amount of calories consumed in a day. More generally, however, the metric should be carefully tailored in a collaboration among health

specialists, behavioural psychologists and computer scientists. The metric should be relevant to the behavioural aspect of interest, but should also be suitable for the design of the underlying machine learning algorithms of the DBCI systems.

Latent Learning without Interfering. Reinforcement learning enhances the knowledge about user's behaviour through exploration. Interacting with the participant through a previously unused tool refines the behavioural model, as we learn more about how the user reacts to this tool. From the practical therapy point of view, however, we face a dilemma: use a tool that is known to elicit a positive behavioural change, or experiment with an unused tool that might yield an even better outcome. In reinforcement learning this dilemma is known as *exploration vs. exploitation trade-off*. Strategies for solving the dilemma in an anticipatory mobile dBCI setting should be aware of the possible irreversible negative consequences of a wrong intervention.

As hinted above, the domain we are working in – human behaviour – is very sensitive, and interventions can lead to serious consequences. This is particularly important when it comes to reinforcement learning that is no more free to explore the action-reaction space rapidly. Rather, the system should try to learn as much as possible about the user without explicitly interacting with the user through behaviour change interventions. Such learning can be realised through *latent learning*, a form of learning where a subject is immersed into an unknown environment or a situation without any rewards or punishments associated to them [25]. Despite the lack of obvious incentives for learning, experiments with both humans and animals show that subjects form a cognitive map of the environment solely because they experience the world around them.

Later, that cognitive map figures in decision making, essentially behaving as a learnt concept. The artificial implementation of latent learning has been demonstrated in [23]. Phones through multimodal sensing can harness latent learning to build a model of the user behaviour with respect to certain actions that correspond to ones targeted by the BCI tools. For example, suppose that in a dBCI that aims to tackle depression the system can provide the user with the suggestion to go out for a dinner with friends. We can get an a priori knowledge of how this suggestion would affect the user, for example if on a separate occasion we detect that the participant went out for a dinner with friends, and we gauge the depression levels, estimated through mobility and physical activity metrics, before and after the dinner. Defining how the expected action – going out with friends – should manifest from the point of view of sensors – e.g., a number of Bluetooth contacts detected, location, time of the day – is one of the prerequisites for practical latent learning. Again, interdisciplinary efforts are crucial to ensure that the detected reaction corresponds to the one that should be elicited by the tool.

HCI in Anticipatory Mobile dBCIs. Effective inclusion of a user in the learning loop requires an appropriate interface between the participant and the system. As noted by Russell et al. [21], a system that autonomously brings decisions and evolves over the course of its lifetime needs to be transparent to the user. Through the user interface such a system must be understandable by the user and capable of review, revision, and alteration. In the behaviour intervention case, interaction should be designed to conform to the best practices in BCI delivery. For example, the content should be framed to emphasise that the tool can help, yet it is fundamental to avoid to harass and patronise the participant.

Large-scale Anticipatory Mobile dBCIs. With billions of smartphones in the world, mobile dBCIs can reach out to an unprecedented number of participants. Context inference through mobile sensing becomes challenging as observable manifestations of a certain behavioural aspect vary over a large number of individuals. Building individual inference models is a slow and cumbersome solution, however, the problem of scale becomes even more pronounced when we talk about tool-effect models, corresponding to the right half of Figure 2. Here, individualism in intervention-reaction space is likely to be highly pronounced – what helps one person quit smoking might not work for a different person. At the same time we want to minimise the search for an intervention tool that works for a specific participant.

In [10] Lane et al. propose community similarity networks (CSN) that enable aggregation and sharing of smartphone sensor data gathered from like-behaved individuals. This joint data significantly speeds up training of machine learning classifiers – similar users can share the same model trained on aggregated data. We propose an extension of CSNs for anticipatory mobile dBCIs. In such networks, a group of similar users can share a behavioural model built on the aggregated tool-effect data. The exact formation of CNSs should be guided by the therapists' expertise, e.g., suggesting that male teenagers respond to a certain intervention in the same way, but also by information learnt through the intervention process via sensing, e.g., realising that participants who walk to work tend to react positively to a certain stress-relief advice. In addition, different intervention tools are likely to have distinct effects in different contexts. By aggregating data over multiple users it might be possible to build context-dependent tool-effect models.

Privacy and Ethics. Location, collocation, physical activity, emotion, together with personal health information can be necessary for mobile dBCIs. Misuse and leaking of such information can be a serious deterrent for mobile dBCI uptake. The most interesting aspect is related to the fact that mechanisms should be put in place in order to guarantee the privacy of data not only describing the current state of the individuals, but also about his/her predicted futures. A more general question arises: who owns the information extracted from the personal data of an individual?

Moreover, there is a need for balancing system performance and privacy in mobile sensing [9]. For example, processing voice data for speaker recognition can be done efficiently if sound samples are sent to a central server for processing, yet, transmission of sensitive speech data presents a privacy risk. In addition, mismatched privacy policies of users who shared the same sensed context can reveal information about people who did not agree to have their information shared, or do not even participate in the dBCI [15]. Finally, the responsibility chain in this domain is yet to be defined. Unsuccessful interventions can have serious consequences. If a depression prevention dBCI fails to recognise a critical episode that leads to the loss of the participant's life, is the blame on the intervention designer, or on the person who devised the underlying machine learning components?

Related Work

Anticipatory computing has recently gained substantial commercial interest. Applications such as MindMeld, GoogleNow (www.google.co.uk/landing/now/), Apple Siri (www.apple.com/ios/siri) and Microsoft Cortana (www.windowsphone.com/en-us/features-8-1) tap into users' personal data in order to provide just-in-time context-

aware services. MindMeld (www.expectlabs.com/mindmeld) uses real-time speech analysis and Web harvesting to enhance online video conferencing with information that the involved parties are likely to find relevant in near future. GoogleNow, Siri and Cortana are more general smartphone applications that provide any information a user may need, but did not explicitly ask for. The implementation details of these systems are not publicly available, however, we can argue that these apps are not anticipatory in the full sense of the word. The applications are context aware and personalised, yet they do not predict nor act in an effort to change the predicted future.

With smartphones behavioural interventions can be tailored to the user and the context in which the user is. Both of these aspects are shown to add value to the intervention [13, 3]. Consequently, the importance of frameworks for practical implementation of smartphone-based context-aware dBCIs has been noted. UBhave [11], for example, enables development of personalised models of user behaviour that are used for deciding on the time and the content of the intervention. However, neither prediction nor anticipation figure in UBhave, and the models are built on past sensed data. Mobile phones represent an attractive platform for psychological research due to the ability to collect ecologically valid data from large populations [14]. In particular, *psychological computing* is advocated by Bao et al. as a class of computing systems that sense user's inner context and utilises it on the core system level [1].

Anticipatory mobile dBCIs rely on sensing to infer user inner state and tightly integrate the user into the mobile computing system, thus fulfil the definition of psychological computing. Moreover, our work is not merely inner context-aware: machine learning models used

for tool-effect evaluation also capture user's reasoning about anticipation.

Summary and Future Work

Ongoing integration of pervasive computing devices with the everyday life, our increased reliance on information originating from mobile phones, and growing ability of such devices to sense, understand, and predict the world around them open up new possibilities for impacting human behaviour in a context-aware, personalised and scalable manner. Current and future context inference, behaviour prediction and action-effect models are necessary parts of anticipatory mobile dBCIs. In addition, such interventions require close integration of psychological tools and machine learning techniques, since the latter are not only used to infer the context that will trigger the former, as is done in traditional mobile dBCIs, but also to model the impact of the tools on the behaviour change. We have presented a possible solution for addressing the machine learning aspect of anticipatory mobile dBCIs, namely through splitting the task into context prediction and tool-effect modelling, reinforcement learning, and latent learning. Yet, we acknowledge that alternative solutions may exist, and are looking forward to future practical implementations of anticipatory mobile dBCIs.

Anticipatory mobile dBCIs are not only relevant for changing human behaviour through predefined intervention tools. Potentially, they also open up a path towards analysing behaviour. These tools may provide guidance, but the actions are in the end selected and performed by the user who bases her decisions on her own judgement. Consequently, the tool-action-effect learning loop includes non-determinism brought in by the user's intuition. While complicating the design of the machine learning components, the non-

determinism effect is actually desirable, since it captures human anticipation decisions in a mathematical model. We believe that such models can prove valuable in behavioural psychology to explain and understand human behaviour.

Finally, we have discussed the design of anticipatory dBCIs around the context inferred from physical built-in sensors on a phone. Digital breadcrumbs that we leave on online social networks, media sharing websites and online portals provide another important source of information, the one that tells a lot about our traits, emotions, intentions. Machine learning tools can turn these breadcrumbs into signals for critical behaviour issues, as evident by the analysis of the Youtube videos of a shooter who killed seven people in Santa Barbara, CA [6]. In future, we plan to investigate how digital breadcrumbs can be mined for information pertaining to anticipatory dBCIs.

Acknowledgements

The authors thank Mihai Nadin from UT Dallas for a thoughtful discussion about anticipatory computing. This work was supported through the EPSRC grant "UBhave: ubiquitous and social computing for positive behaviour change" (EP/I032673/1).

References

- [1] Bao, X., Gowda, M., Mahajan, R., and Choudhury, R. R. The Case for Psychological Computing. In *HotMobile'13* (Jekyll Island, GA, USA, February 2013).
- [2] Burbey, I., and Martin, T. L. A Survey on Predicting Personal Mobility. *International Journal of Pervasive Computing and Communications* 8, 1 (2012), 5–22.
- [3] Burns, M. N., Begale, M., Duffecy, J., Gergle, D., Karr, C. J., Giangrande, E., and Mohr, D. C. Harnessing Context Sensing to Develop a Mobile Intervention for Depression. *Journal of medical Internet research* 13, 3 (2011).
- [4] Campbell, A., and Choudhury, T. From Smart to Cognitive Phones. *IEEE Pervasive Computing* 11, 3 (2012), 7–11.
- [5] Cole-Lewis, H., and Kershaw, T. Text Messaging as a Tool for

- Behavior Change in Disease Prevention and Management. *Epidemiologic Reviews* 32, 1 (2010), 56–69.
- [6] Foschini, L. Could Machine Learning Have Helped Prevent the Isla Vista Shooting?, May 2014. <http://lucafoschini.com/2014/05/25/machine-learning-isla-vista-shooting.html>.
- [7] Klasnja, P., and Pratt, W. HealthSense in the Pocket: Mapping the Space of Mobile-Phone Health Interventions. *Journal of biomedical informatics* 45, 1 (2012), 184–198.
- [8] Lane, N. D., Choudhury, T., Campbell, A., Mohammad, M., Lin, M., Yang, X., Doryab, A., Lu, H., Ali, S., and Berke, E. BeWell: A Smartphone Application to Monitor, Model and Promote Wellbeing. In *Pervasive Health'11* (Dublin, Ireland, May 2011).
- [9] Lane, N. D., Miluzzo, E., Lu, H., Peebles, D., Choudhury, T., and Campbell, A. T. A Survey of Mobile Phone Sensing. *IEEE Communications Magazine* 48, 9 (2010), 140–150.
- [10] Lane, N. D., Xu, Y., Lu, H., Hu, S., Choudhury, T., Campbell, A. T., and Zhao, F. Enabling Large-scale Human Activity Inference on Smartphones using Community Similarity Networks (CSN). In *UbiComp'11* (Beijing, China, September 2011).
- [11] Lathia, N., Pejovic, V., Rachuri, K., Mascolo, C., Musolesi, M., and Rentfrow, P. J. Smartphones for Large-Scale Behaviour Change Intervention. *IEEE Pervasive Computing* 12, 3 (2013).
- [12] Lu, H., Mashfiqui Rabbi, G. T. C., Frauendorfer, D., Mast, M. S., Campbell, A. T., Gatica-Perez, D., and Choudhury, T. StressSense: Detecting Stress in Unconstrained Acoustic Environments using Smartphones. In *UbiComp'12* (Pittsburgh, PA, USA, September 2012).
- [13] Lustria, M. L. A., Noar, S. M., Cortese, J., Stee, S. K. V., Glueckauf, R. L., and Lee, J. A Meta-Analysis of Web-Delivered Tailored Health Behavior Change Interventions. *Journal of Health Communication* 18, 11 (2013), 1039–1069.
- [14] Miller, G. The Smartphone Psychology Manifesto. *Perspectives on Psychological Science* 7, 3 (2012), 221–237.
- [15] Pejovic, V., and Musolesi, M. Anticipatory Mobile Computing: A Survey of the State of the Art and Research Challenges. Submitted for Publication. Preprint available on ArXiv: arXiv:1306.2356, June 2013.
- [16] Pejovic, V., and Musolesi, M. InterruptMe: Designing Intelligent Prompting Mechanisms for Pervasive Applications. In *UbiComp'14* (Seattle, WA, USA, September 2014).
- [17] Rachuri, K. K., Mascolo, C., Musolesi, M., and Rentfrow, P. J. SociableSense: Exploring the Trade-offs of Adaptive Sampling and Computation Offloading for Social Sensing. In *MobiCom'11* (Las Vegas, NV, USA, September 2011).
- [18] Rachuri, K. K., Musolesi, M., Mascolo, C., Rentfrow, P. J., Longworth, C., and Aucinas, A. EmotionSense: A Mobile Phones based Adaptive Platform for Experimental Social Psychology Research. In *UbiComp'10* (Copenhagen, Denmark, September 2010).
- [19] Rosen, R. *Anticipatory Systems*. Pergamon Press, Oxford, UK, 1985.
- [20] Roshanaei-Moghaddam, B., Katon, W. J., and Russo, J. The Longitudinal Effects of Depression on Physical Activity. *General Hospital Psychiatry* 31, 4 (2009), 306–315.
- [21] Russell, D. M., Maglio, P. P., Dordick, R., and Neti, C. Dealing with Ghosts: Managing the User Experience of Autonomic Computing. *IBM Systems Journal* 42 (2003), 177–188.
- [22] Scellato, S., Musolesi, M., Mascolo, C., Latora, V., and Campbell, A. T. Nextplace: A spatio-temporal prediction framework for pervasive systems. In *Pervasive'11* (San Francisco, CA, USA, June 2011).
- [23] Stolzmann, W., and Butz, M. *Learning Classifier Systems, From Foundations to Applications*. Springer-Verlag, Berlin, Germany, 2000, ch. Latent Learning and Action Planning in Robots with Anticipatory Classifier System, 301–321.
- [24] Sutton, R. S., and Barto, A. G. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, USA, 1998.
- [25] Tavris, C., and Wade, C. *Psychology in Perspective*. Longman, New York, NY, USA, 1997.
- [26] Wang, T., Cardone, G., Corradi, A., Torresani, L., and Campbell, A. T. WalkSafe: A Pedestrian Safety App for Mobile Phone Users Who Walk and Talk While Crossing Roads. In *HotMobile'12* (San Diego, CA, USA, February 2012).