Challenges for Designing Notifications for Affective Computing Systems

Mary Czerwinski
Microsoft Research
One Microsoft Way
Redmond, WA, 98052 USA
marycz@microsoft.edu

Ran Gilad-Bachrach
Microsoft Research
One Microsoft Way
Redmond, WA, 98052 USA
rang@microsoft.edu

Shamsi Iqbal
Microsoft Research
One Microsoft Way
Redmond, WA, 98052 USA
shamsi@microsoft.edu

Gloria Mark
Donald Bren School of Information and Computer Science
University of California, Irvine
Irvine, CA 92617
gmark@uci.edu

Ubicomp/ISWC'16 Adjunct , September 12-16, 2016, Heidelberg, Germany
© 2016 ACM. ISBN 978-1-4503-4462-3/16/09…$15.00
DOI: http://dx.doi.org/10.1145/2968219.2968548

Abstract
Affective computing systems blend technical and social elements and present challenges for designing notifications. As the number of applications expand that utilize notifications, we need to consider that competing for attention is a concern. With respect to promoting positive health behaviors, notifications need to be adapted to times when a user is vulnerable and receptive for a health intervention but also to users' attentional states. In this paper we outline affective computing projects and in the workshop will discuss challenges of designing interventions to promote positive behaviors in real world environments.

Author Keywords
Notifications, Affective computing, just-in-time interventions, cognitive overload, wearables

ACM Classification Keywords
H.5.3 [Information Interfaces and Presentation (e.g., HCI)]: Group and Organization Interfaces; K.4.m [Computers and Society]: Miscellaneous.

Introduction
Affective computing systems blend technical and social elements. For example, determining a user's emotional state may be accomplished by prediction algorithms that run on sensor data (e.g., from wearables like

Microsoft Band, or from a mobile phone), external or environmental data (e.g., weather), data reported directly by the user (e.g., through lifelogging), data inferred from the user (e.g., from queries with a personal digital assistant), and/or data reported by people close to the user (e.g., a parent).

Once a potentially problematic emotional state has been determined by a computer system– for example, the user who is trying to lose weight is predicted to be about to eat – then the most effective intervention may rely on a combination of technical and social approaches – for example, whether an intervention is needed, when to present it and in what form. A machine learning algorithm may be used to predict whether an intervention is warranted for the particular moment: for example, if it is breakfast time and the user has not eaten yet today – or the user is about to consume food that is part of a healthy diet, an intervention that discourages eating at that moment might be counterproductive. When interventions are indeed warranted, they may also take various socio-technical forms, such as an auto-generated reminder, or a text message composed by a concerned friend. But the user's context is extremely important as to the style and format of that intervention to decrease its disruptive nature and to increase is acceptability by the user.

Affective computing applications using notifications

As these examples illustrate, there are many possibilities for both the sensing and intervention components of affective computing. The research in this area consists primarily of developing prototype technologies or approaches for each component, and carrying out studies of their use by research subjects, to determine the most effective components, contexts and temporal requirements. Affective computing interventions can be generalizable to many health-related scenarios of personal and societal importance, including chronic disease prevention and management, smoking cessation, weight control, excessive drinking, anger management, and many other problematic, often widespread behaviors.

In the near term, affective computing systems with adaptive, just-in-time interventions could be directly adopted by consumers seeking to change their own behaviors. They could also be adopted by physicians and psychologists for patient studies addressing medical and/or mental health issues. Such just-in-time interventions are a treatment design in which intervention options are delivered in real-time, and are adapted to address the immediate and changing needs of individuals in daily life. We present some examples below.

Just-in-time eating prediction: Being able to detect when a poor behavioral decision is about to be made can enable more effective interventions than reacting afterward. MSR's affective computing team has developed a novel wearable sensing framework [7] that predicts moments when a user is *about to* eat, enabling intervention when it is most likely to be effective. The framework consists of a multimodal array of sensors combined with signal processing, machine learning and personalization algorithms. This prototype framework enabled training of an "About-to- Eat" moment classifier with an average recall of 85%. The sensors involved capture physical activity, location, heart rate, electrodermal activity, skin temperature and caloric expenditure. For the experimental prototype, these sensors came from a combination of a Microsoft Band 1.0, an Affectiva Q sensor, a wearable microphone, and an Android phone app.

Vetted interventions: MSR surveys on mTurk have identified about 100 micro-interventions that people want and would actually do (e.g., 10 seconds of deep breathing in response to stress), while ruling out others recommended in the psychology literature (e.g., journaling). Other interventions that have shown promise include scribbling and gamified calming (e.g., pushing dots back and forth 20 times in an app). Again, the user context will need to be determined to understand which intervention and when the user would be willing to use one of these interventions. Also, the user's device will be extremely important to consider as well.

Getting people to change their eating and exercise habits is so difficult that even small changes are considered big victories. Our novel approaches could achieve those kinds of victories. For example, the just-in-time eating prediction work [7] just won a best paper award, demonstrating its novelty. Adding "about to eat" prediction capability into MS Health would give us a unique tool for promoting healthy eating habits.

Modeling good behavior. Another application of just-in-time interventions, ParentGuardian, provided support to parents of children with ADHD during challenging moments, and enabled them to model good behavior while "under fire." The MSR study of ParentGuardian with 16 Microsoft children ended two years ago, but half of the children's parents continued using the application 2 years later.   This study used Affectiva's Q sensor for its electrodermal sensing capability to measure parental stress but the parents didn't have the sensors after the 2-week experiment, which suggests that the interventions in the app were useful reminders for modeling good behavior with their children.

We expect to be able to improve the accuracy of the "just in time eating intervention" work through the use of more contextuaizedl and sensor data. More generally, this work could be generalized into domains beyond diet control, as noted above.

Taking work breaks. Another application for just-in-time interventions is in the workplace, in detecting when it is time for workers to take a break. Studies show that work breaks can serve to refresh people, and can lead to higher productivity, greater satisfaction, and higher motivation, [cf 2]. In previous work by our team at MSR, using an array of sensors along with experience sampling, we were able to track how attention focus varied over the workday with information workers. We discovered that there were identifiable temporal patterns throughout the workday: rhythms where attention is focused, where work is experienced as routine or rote, and where the work experience is that of boredom [5]. Notifications can serve to remind people to take a break from work, especially in times when people feel bored or stressed. In fact, recent research has shown that leaving the office and walking in a nature setting can serve to significantly enhance creativity when returning to the workplace [4].

Challenges for notification designs for affective computing systems

So far, we have identified several valuable uses for notifications as just-in-time interventions in improving quality of life. We now identify some challenges for designing effective just-in-time interventions with respect to how they impose demands on attention.

Competition for attention. As we expand our applications for uses of notifications, we need to consider that it will become more and more competitive to attract attention from notifications. Competing for attention will become an increasing concern as we move towards an IoT environment where more and more devices in varied contexts will compete for our attention. In order to act on just-in-time interventions, people need to attend to the notifications. We need to develop effective notification systems that subtly nudge users towards the desired behavior without overburdening the user or appearing disruptive. We thus need to ask: how do we time and contextualize notifications so as to optimize for minimal disruption of the user's attention while still promoting good eating habits, healthy behaviors or reducing stress in the workplace. We need to balance the costs and benefits of disruption vs. the potential health improvements from receiving them.

To address this challenge, we can draw on some research to inform us. Following the resource usage principle of threaded cognition [8, or multiple resource theory (Wickens, 89ish)], once cognitive resources become available from a primary task, they can be used for another task. How this relates to notifications concerns the modality of notifications. If a person is doing a visual task, e.g., watching a video, they will be less likely to notice a notification that uses visual media, e.g., a blinking icon, as the resources are shared, i.e., both the primary task and the notification use visual resources. Thus, notifications that use a modality different from the task-at-hand, will more likely be attended to.

When to interrupt with notifications. A challenge for deploying notifications is finding the right time to intervene. Work by Iqbal and Bailey [3] found that interruptions that occur at low workload moments have lower negative effect on task performance and affective states such as frustration. When the cognitive workload of a task is high, e.g., in the middle of writing a paragraph, an interruption is more disruptive than when the workload of a task is lower, e.g., when one has come to a breakpoint in the task, such as finishing a section of writing.

To address this challenge, we should consider that notifications that occur at low workload moments will likely be less disruptive, and serve a greater chance of being acted upon. However, the notion of workload needs to be considered alongside the timeliness and importance of the notification in order to retain the effectiveness of the notification.

Conclusions

Thus, the design of just-in-time interventions needs to juggle several factors. First, they need to identify moments when the intervention can be effective. The adaptive, just-in-time intervention approach is the precise provision of timely support whenever the person is (a) open to positive changes, and (b) receptive [6]. Just-in-time interventions need to provide support based on ongoing assessments of the individual's state and ecological context. The goal is to offer the right type of support precisely when the person is in a state of vulnerability, opportunity and receptivity.

Second, notifications need to be adapted to users' attentional states. A user may be vulnerable and receptive for a health intervention, but they may not be in a cognitive state where they can effectively process the notification. If a person does not have sufficient cognitive resources to process the notification, they may not respond. Techniques such as pupil dilation have been shown to be a promising indicator of mental workload [e.g., 1], i.e., moments when a person may be receptive to attending to a notification. However, more research is needed in the ubiquitous computing field to detect and understand comprehensive views of users in terms of their affective and cognitive states, as they are situated in an environmental context, in order for just-in-time interventions to be effective.

Limitations

There are several important limitations that we need to be aware of when designing just-in-time interventions. We have outlined challenges in deploying interventions, which are informed by the collection of sensor data. However, collecting sensor data poses its own challenges: it is important to collect data as unobtrusively as possible so as not to affect the user's state, e.g., to not impose additional stress which could affect the effectiveness of technology deployment. Another potential limitation is in not collecting data over a long enough period of time. Longitudinal studies enable more generalizability of the models. Lastly, another possible limitation is in not using a large enough sample size for predictive purposes. A larger sample can improve the personalization of the models learned.

We will present our projects described in this position paper at the workshop and will discuss the results and challenges of designing interventions to promote positive behaviors in real world environments.

References

1. Jackson Beatty and Brennis Lucero-Wagoner, 2000. The pupillary system. *Handbook of psychophysiology*, *2*, pp.142-162.
2. Charlotte Fritz, Allison M. Ellis, Caitlin A. Demsky, Bing C. Lin, and Frankie Guros. 2013. Embracing work breaks: Recovering from work stress. Organizational Dynamics 42, 4: 274-280.
3. Shamsi T. Iqbal and Brian P. Bailey, 2005, April. Investigating the effectiveness of mental workload as a predictor of opportune moments for interruption. In*CHI'05 extended abstracts on Human factors in computing systems* (pp. 1489-1492). ACM.
4. Marily Oppezzo and Daniel L Schwartz. 2014. Give your ideas some legs: The positive effect of walking on creative thinking. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 40, 4 (2014), 1142–1152.
5. Gloria Mark, Shamsi T. Iqbal, Mary Czerwinski, and Paul Johns (2014). Bored Mondays and focused afternoons: The rhythm of attention and online activity in the workplace. *Proceeding of the thirty-second annual SIGCHI conference on Human factors in computing systems (CHI'14), ACM Press.*
6. Inbal Nahum-Shani, Shawna N. Smith, Ambuj Tewari, Katie Witkiewitz, Linda M. Collins, Bonnie Spring, and S. Murphy. "Just in time adaptive interventions (jitais): An organizing framework for

ongoing health behavior support."*Methodology Center technical report* 14-126 (2014).

7. Tauhidur Rahman, Mary Czwerwinski, Ran Gilad-Bachrach, and Paul Johns, 2016. Predicting “About-to-Eat” Moments for Just-in-Time Eating Intervention. *Proceeding of ACM Conference on Digital Health 2016 (DH'16),* ACM Press.

8. Dario Salvucci and Niels Taatgen, 2010. *The multitasking mind*. Oxford University Press.