

BioCrystal: An Ambient Tool for Emotion and Communication

Asta Roseway¹; Yuliya Lutchyn¹; Paul Johns¹; Elizabeth Mynatt²; Mary Czerwinski¹

¹Microsoft Research, One Microsoft Way, Redmond, USA; ²Georgia Institute of Technology, Atlanta, USA
{astar, v-yulutc, PaulJohns@microsoft.com, mynatt@cc.gatech.edu, marycz@microsoft.com}



ABSTRACT

In this paper we present the BioCrystal – a biofeedback device that uses physiological data to evaluate user's affective states in real-time and signals the states via an ambient display. We evaluated the BioCrystal during a 2-week, *in situ* multi-method study during which ten users collected over 115 hours of usable data. Users' comments suggested high utility of such a biofeedback device for self-awareness, stress-management and interpersonal communication. Quantitative data confirmed that the BioCrystal met the criteria of an ambient display, and significantly improved users' ability to control their stress. We discuss practical applications and suggest directions for future development.

Keywords

Moods, Emotions, Ambient Display, Biofeedback, Affective Computing, Workplace Communication, Field Study.

INTRODUCTION

Emotions are a critical part of human experience, and have a significant impact on all stages of human perception, cognition, and behavior. Emotions directly influence how we perform everyday tasks, interact with each other, learn, work, and even make important decisions (Loewenstein & Lerner, 2003). Emotions, especially negative ones, also affect our physical and mental well-being. With rapidly evolving methods for measuring affective states, it is not surprising that understanding, communicating, and even manipulating emotions are now among the central research topics in the domain of human-computer interaction. Our study is an attempt to contribute to this growing body of research by designing a novel, personal technology that would monitor and signal users' affective states, enable them to selectively communicate this information to others, and, hopefully, help to manage their affect.

Over the past decade, researchers have employed a variety of methods to assess affect: physiological measurements (AlZoubi, D'Mello, & Calvo, 2012; Cacioppo & Tassinary, 1990), recognition of facial expressions (Stefani, Mahale, Pross, & Bues, 2011), speech analysis (D'Mello, Craig, Witherspoon, McDaniel, & Graesser, 2008), linguistic analysis of text and status updates in social media (Golder & Macy, 2011), and, more recently, users' self-reports through mobile applications (Church, Hoggan, & Oliver, 2010). Each of these approaches has its benefits, but also suffers from one or more shortcomings: poor accuracy, low external validity, privacy issues, intrusiveness, limited usability and an inability to share the information (Geven, Tscheligi, & Noldus, 2009). We aim to address some of these challenges by synthesizing methods and findings accumulated through previous research on affect-

awareness. Specifically, the goal of our study was to create a mood-management device that would meet the following criteria: 1) it would assess users' affect via objective physiological data, 2) data would be collected in situ over a reasonably long period of time, 3) the device would provide real-time feedback, 4) the device would have an ambient display (i.e., it would be unobtrusive and not require a user's full attention to use), 5) the feedback information would be intuitive and easy to read (i.e., would not have a complex meaning map); 6) the user would have an option to share information about affect if desired, and, finally, 7) the device would demonstrate measurable stress-management benefits.

RELATED WORK

Although there is a growing body of research on affect-sensing, the research on individual mood-monitoring technology is still relatively scarce, and there is even less information on how individuals evaluate and use their affect-sensing devices in real life.

One of the first successful attempts to create an affect-recognition system was made by Kapoor and Picard (2005), who developed a multi-sensor system predicting students' interest in a learning environment. Researchers used facial expressions and body posture to predict user's state, and were able to achieve high accuracy of affect recognition (86%). A more recent example is work by Stefani and colleagues (Stefani et al., 2011) who created SmartHeliocity – technology that uses facial expressions to determine human emotions and controls a color-changing bulb in order to influence emotional state of the user in the workplace: e.g., turns red to energize a tired or sleepy employee, or turns green to calm a stressed one. Another notable example is AffectAura – an "emotional prosthetic" that allows users to reflect on their emotional states over long periods of time (McDuff, Karlson, Kapoor, Roseway, & Czerwinski, 2012). AffectAura used a system of sensors to collect continuous physiological, audio, video and contextual data to assess individual's affective state and visualize it for user reflection. To the best of our knowledge, this is the only study on individual emotion-sensing technology that formally tested utility of the device in situ and was able to demonstrate its functional benefits. Building on the work of the AffectAura designers, we also aim to create an affect-measuring device that will collect continuous objective data and provide users with the visual representation of their affect. However, our solution is different in several critical ways. First, our affect-prediction engine will run only on physiological data collected from small wearable sensors. Of course, as demonstrated by McDuff and colleagues (McDuff et al., 2012), using different channels of information and a range of sensors allows researchers to capture more diverse information and increases accuracy of affect classification. However, we felt that such a multimodal sensor set-up makes the device more intrusive and limits its usability. Excluding video, audio and activity log data from the prediction model significantly expands the number of usage scenarios to health care and educational settings, shopping malls, homes etc.

Another critical distinction from the AffectAura study is our approach to visualizing data. Visualization of the user's affective state has been shown to increase a user's self-awareness and promote reflection on one's own internal state (Liu, 2004). Therefore, representation of data is as important as accuracy of the prediction engine, and will be critical to the quality of the user's experience with the affect-monitor. In the present study we wanted to create an ambient display that would require only minimal users' efforts to access and process affect information. Ambient displays are designed to be minimally attended from outside of a person's direct focus of attention, providing a level of pre-attentive processing without being unnecessarily distracting. Such technologies are designed to blend smoothly in the environment and provide an information channel that is "glanceable" yet can be easily ignored (Hazlewood, Stolterman, & Connelly, 2011). All these properties of ambient displays appear particularly relevant for the technology with stress-management properties.

Finally, we wanted the users of our device to have an option of selectively sharing their affect information, thus making the device useful for interpersonal communication. The majority of existing solutions for detection and presentation of mood information do not facilitate the sharing of information with others. A few designs that do support a sharing feature, imply high involvement of both the user and the audience – e.g., sharing affect via a phone app in a Twitter-like manner (Church et al., 2010; Gay, Pollak, Adams, & Leonard, 2011; see also MacLean, Roseway, & Czerwinski, 2013 for a notable exception). Our goal was to design a technology that would allow its user to communicate affect in real time in an unobtrusive manner.

PREDICTING AFFECT

We defined users' affective state according to the Circumplex Model of Affect, which posits that all affective states can be mapped onto a two-dimensional system formed by the most fundamental properties of the affective experience: valence (negative – positive affect) and arousal (calm – energetic) axes (Russell, 1980). In developing a tool for visualizing users' affective state, an important first task was to create an affect-prediction engine: to select physiological parameters for capturing user affect, set-up appropriate sensors and develop an affect recognition model based on sensed data. Affective computing studies employ many different modalities and features to predict mood or emotions (AlZoubi et al., 2012; Calvo & D'Mello, 2010; Picard, Vyzas, & Healey, 2001). We chose to estimate arousal based on the user's electrodermal activity (EDA). EDA is widely considered one of the strongest features that can be used for detecting arousal (Boucsein, 2012). The signal tends to gradually increase with arousing emotions and decreases for deactivating ones. We measured EDA via an Affectiva Q Sensor - a cloth wristband with a built-in sensor (Fig. 1). Valence was determined based on the participants' heart rate variability, using the input signal from a Zephyr HXM ElectroCardioGram monitor - a chest band sensor. Both devices communicated wirelessly with the computer. Previous research has shown that wearable sensors like these are an effective means for capturing *in situ* workplace affect, as they enable the continual capture of affect with minimal disruption for the participant (Mark, Iqbal, Czerwinski, & Johns, 2014, February).

The user's position on the arousal axis (anchored at -100= lowest arousal and +100=highest arousal) was predicted by comparing the user's EDA over the prior 10 seconds to the user's average, minimum, and maximum EDA values over an entire session. The valence axis, also anchored at -100/+100, was determined by comparing the heart rate variability for the user over the prior 5 minutes to the variability over their history (multiple sessions). Because the Affectiva EDA readings can vary depending on where the band is situated on the wrist and how tightly the strap is positioned, we did not compare the current EDA to the user's lifetime history of EDA readings (as opposed to the ECG monitor, where the historical comparison is valid). The HRV metric used in this study was the root mean sum of the squared differences (rMSSD), calculated using a rolling five minute window.



Fig. 1. Affectiva Q Sensor (left) and Zephyr HXM ElectroCardioGram monitor (right)

BIOCRYSTAL DESING

To ensure that the interface meets the criterion of an ambient device and creates a compelling experience for the user, we started its development with an informal user study and an extensive iterative design exploration. We recruited eleven volunteers (9 female, all knowledge workers at a major corporation) who provided us with user insight throughout the design process. We described the concept of the affect monitor to our participants and asked them to comment on the idea, with special focus on users' stress-regulation habits and perceived value of the device. Series of formative semi-structured interviews confirmed our intuitive predictions. Participants agreed that the device would be highly useful for monitoring general and work-related stress, and expressed their willingness to use such a device. Participants also noted that the affect display would be a good aid for interpersonal communication: when displayed publically, it could signal when its user is stressed, focused or calm, and, therefore, can or cannot be interrupted. At the same time, users also had a concern that visibility of their personal affective information could interfere with their efforts to display only "professional" emotions in the workplace. A solution to this conflicting demands would be an interface with "private" and "public" display modes. These and other secondary insights directly influenced the first iterations of the interface design. It was decided that the affective information should be communicated by the color-changing light. This is a very popular solution in ambient devices because lights are "glanceable", easy to embed in many everyday objects, and fit in virtually any environment (Chang, Resner,

Commented [YL1]: We must provide signed permission forms for all copyrighted images used in the paper. I took these two pictures from manufacturers' websites. I could try to get permissions to use these photos, but it would be WAAY faster and easier to include our own pics of the sensors. Do we have any?

Koerner, Wang, & Ishii, 2001; Olivera, Rivas, & Iturriaga, 2013; Rea, Young, & Irani, 2012; Stefani et al., 2011). We created 20 versions of the ambient interface with an embedded color-changing LED light: Zen rock garden, little Buddah, Star War lamp set, hanging lanterns, fabric flower lamp, egg crystal, lotus crystal, raw (asymmetrical) crystal, and others. We featured all prototypes in a mockup catalogue, and asked our participants to rank them (Fig. 2).




Fig. 2. Page from the mockup catalogue.

Participants chose the Crystal as their most favorite design, so it was selected as an interface for our affect-sensing device. The BioCrystal consists of a small platform base with an LED light and a small, plastic, detachable sphere-shaped shell (Figure 3). Each crystal was driven by an Arduino Uno and was tethered via USB to a laptop to ensure quality connection and consistent power. The shell casings were designed in CAD and 3D printed on a MakerBot Replicator 2 using transparent PLA and basic black PLA. In the main usability study, each user received two shell crystals to be used interchangeably: a transparent, white one for the public display of their mood, and a black, opaque one for privacy mode. The BioCrystal displayed a user's mood based on combined input from two sensors described above, and reported it in real-time via an LED light that changes its color from white, to red, green, dark blue, or light blue (Fig. 4).





Figure 3. The BioCrystal with an opaque (black) and a transparent (white) shells.

Figure 4. Color-affect map

As can be seen in the color-affect map, areas of white color were “neutral” affect zones, where the Crystal displayed no saturation. We did not want to complicate the study with a dense set of highly abstract mappings of colors to meanings, or use anything that would require significant training or effort in order to be understood. Previous research and our own pre-test confirmed that these color-affect associations were easy to understand for the US participants (Madden, Hewett, & Roth, 2000; Stefani et al., 2011).

USER STUDY

To understand how individuals use the BioCrystal, to evaluate its utility, and to determine directions for further improvement, we conducted a usability study that aimed to answer several main questions:

- RQ1. What is the users' overall evaluation of the Crystal?
- RQ2. Does the Crystal meet criteria of an ambient device?
- RQ3. What users see as the main benefits and shortcomings of the Crystal?
- RQ4. How do people use the Crystal? How exactly do they interact with the device?
- RQ5. How accurate is the Crystal relative to user's perception of their own affective state?
- RQ6. What are the benefits and risks of using of the Crystal for interpersonal communication?
- RQ7. Do personality and situational factors influence usage and evaluation of the Crystal?

Participants

Participants were volunteers working in a research division of a large corporation. Work setting is one of the most relevant contexts for exploration and adoption of the affect-monitoring device: workplace is where most people spend significant portions of their time, and where affect has significant impact on both process and outcome of users' daily activities (see Brief & Weiss, 2002, for review). Multiple studies demonstrated that emotional states of the employees' have significant influence on their decision making (Forgas, 1995), creativity (Baas, De Dreu, & Nijstad, 2008), negotiation strategies (Pillutla & Murnighan, 1996), and job performance (Croppanzano, James, & Konovsky, 1993).

Ten people (nine females, age range of 25-60 years old) completed the study. All participants were knowledge workers and had a variety of job roles: general administrative support, event management, support for senior executives, and engineering support for product development and testing. No participants were color-blind. Participants were compensated with a \$250 gift card.

Procedure

Before starting the study, all participants filled-out a pre-study survey with demographic information and a color-blindness screening question (see Table 1 for a detailed description of all self-reported measures). The study was conducted in two steps. During the first five days (“week 1”), we collected physiological and self-report data in order to establish a baseline and build an affect-prediction model for each participant. At the beginning of the study we meet with each participant individually to explain the study procedure, install the software, and instruct on how to use the sensors. Participants collected physiological data during regular work activities, either at their own desk, or away while using a laptop. Participants were asked to work as they normally would throughout the workday. They were instructed to collect at least three hours of data each day, but decided themselves when to run a session of data collection. Participants could stop and resume data collection as necessary. In addition, participants also replied to several experience sampling probes that were programmed to popup on the computer monitor several times during each data collection session. At the end of each study day they also filled out an end-of-day survey reporting their mood and describing the work day. No affective feedback was provided to the participants during the first week.

The next phase (“week 2”) was identical to the first one, except this time the participants had their own BioCrystal installed on their desktops. We asked participants to use the BioCrystal for a minimum of 5 days. This was done to ensure that the Crystal overcame any initial novelty effects and became fully integrated with their daily routines (Hazlewood et al., 2011). The participants could place the BioCrystal anywhere on the desk, and were told that they

Commented [MC2]: Do you mean on top of their desks? Sound like you are referring to the virtual desktop.

could start and finish using the BioCrystal as they wanted. The end-of-day survey was edited to include BioCrystal evaluation questions. After completing the study, participants filled out a post-study survey with personality inventory and took part in a semi-structured interview in which they expressed their overall impression of the Crystal and commented on specific benefits and shortcomings of the device.

Measures

Hazlewood and colleagues (2011) pointed out that measuring effectiveness of ambient displays is a non-trivial task, because the very notions of *user* and *use* can be somewhat ambiguous. For example, are we using the device just by being in the same space and processing it at the pre-attentive level, or does it require direct observation and consideration? What are the expected outcomes of the device use? In case of the affect-monitoring ambient display, such as the BioCrystal, we hope that the ultimate outcome is user's improved emotional well-being, including improved stress control. But what are the measures of success, precisely? Possible outcomes range from mere affect awareness to affect control and even higher-level cognitive and behavioral changes, such as improved productivity, workplace communication and job satisfaction. In attempt to capture these and other possible benefits of using the Crystal, we employed a variety of open-ended and structured measures (Table 1 provides a summary of the assessment tools used in the study).

Affective state. Physiological indicators of the affective state were complemented with self-reports collected via experience sampling throughout the day. Experience sampling was done with custom built software that presented a probe – a small pop-up window on the computer screen that presented participants with the 2x2 affect-arousal grid and prompted them to click with their cursor on the point on the grid that best expressed their feeling "right now" (see Figure X). Experience sampling like this pairs well with automatically logged data, and has been successfully used for measuring mood (Hektner, Schmidt, & Csikszentmihalyi, 2007; Mark, Iqbal, Czerwinski, & Johns, 2014, April).

Commented [YL3]: Can we make a screenshot of this popup window?

The popup window also measured participant's current level of engagement and challenge. These cognitive dimensions of the user state are known to be highly relevant in capturing user's experience at the workplace, and were found to be tied to the users' affective state (Mark, Iqbal, Czerwinski, & Johns, 2014, April). Engagement is the degree to which people feel involved in or distracted from their work, and is particularly relevant in information work (Mark, Iqbal, Czerwinski, & Johns, 2014, February). Challenge refers to the task complexity and has long been emphasized for its significant contribution to the experience of "flow" and overall job satisfaction (e.g., Herzberg, 1993). As we learnt in our informal pre-study, communicating states of engagement and challenge were highly desirable interpersonal communication functions of an ambient display.

Crystal evaluation. On week 2, the end-of-day survey included a set of items that measured users' attitude toward the Crystal. On a 5-point scale, participants rated accuracy, intrusiveness, usefulness, ease of use of their device and indicated how much they were paying attention to it. User's also had a chance to comment on their experience in an open-ended format. To capture any indirect consequences of using the BioCrystal, the daily survey also asked participants a series of questions about the flow and the results of their work that day. Within two days of completing the study, all participants were interviewed about their experience with the BioCrystal. The semi-structures interview lasted 15-30 minutes, and probed for the following issues: the overall attitude toward the BioCrystal, intrusiveness, accuracy, color mapping, effect of the device on the work routine, mood and mood management, and interpersonal communication effects.

Personality variables. Previous research suggests a significant relationship between some personality traits, affective states, affect-management strategies and reactions to emotional regulation demands that exist in the workplace (e.g., Bono & Vey, 2007). To explore the role of individual differences on the use of our technology, we included measures of the key dispositional variables most relevant to our scenario: emotional intelligence, the Big Five personality inventory, and a self-monitoring questionnaire. Emotional intelligence is the ability to accurately perceive and understand one's own and other people's emotions, to control one's own affective state, and to use emotional information to guide thinking and behavior (Salovey & Mayer, 1990). People high on EI are usually also more successful in emotional regulation. The Big Five inventory measures five well-validated key dimensions of personality: agreeableness, conscientiousness, extraversion, openness and neuroticism (John & Srivastava, 1999). Studies have found relationships between some of these traits and affect: e.g., neuroticism is generally associated with negative affective states, and extraversion is associated with positive affective states (Watson & Clark, 1992).

Mark et al. also found that conscientiousness is positively related to the engagement rating (2014, February). Self-monitoring refers to the tendency to monitor one's behavior and adapt it according to the demands or expectations of a particular situation (Gangestad & Snyder, 2000; Snyder, 1974). High self-monitors are generally very effective in emotion regulation – not only are they more accustomed to managing emotions, but they also use a distinct strategy for that: Bono and colleagues found that in order to manage their emotions, high self-monitors engage in deep acting, which allows them to produce the required emotion without experiencing increased stress or appearing insincere (Bono & Vey, 2007).

Interpersonal effects of the BioCrystal use. A social-functional approach to emotion holds that emotions not only influence those who experience them, but also those who observe them (Frijda & Mesquita, 1994; Van Kleef, Homan, & Cheshin, 2012). For example, people may “catch” other people’s emotions (laughter), may experience complementary emotions (sadness often evokes sympathy, anger evokes fear, etc.); people also make inferences about individuals and situations based on the affective states they observe in others. Therefore, using an affect-sensing technology with a scrutable interface can have not only individual, but also interpersonal ramifications. This aspect of the device usage has also emerged as one of the main themes in our informal pre-test.

On one hand, visualizing one’s emotions and displaying them via an external device can aid interpersonal communication by making affective states clearer to the observer. Indeed, only a few basic emotions have easily recognizable expressions (e.g., disgust, surprise, excitement, anger), and these expressions – as well as emotions themselves – are very fleeting (Ekman & Friesen, 2003). Most of the time, our moods are in the low-arousal zone and are not easily recognizable. By clearly labeling user’s affective state, the Crystal could facilitate communication in the workplace. On the other hand, this very ability of the Crystal may interfere with users’ efforts to convey only professional emotions, i.e., masking negative emotions, making sure emotions are neither excessive nor poorly timed, and other efforts to express emotions strategically (Hochschild, 1979; Kramer & Hess, 2002). In other words, the BioCrystal can make it difficult for users to apply professional rules for their emotion display. We included a number of measures to explore both desirable and undesirable interpersonal effects of the BioCrystal.

Table 1.

Summary of self-report measures used in the study.

Concept	Instrument	Sample item(s)	Time administered
Color-blindness (screener)	Multiple-choice	Which of the following statements best describes your ability to differentiate between colors? a) I can see all colors and hues (questions continue); b) I am totally or partially color blind; c) I am not sure	Pre-study survey
Age	Multiple-choice	25 or less,...,60 or more	Pre-study survey
Gender	Multiple-choice	Male; female	Pre-study survey
Education	Multiple-choice	High school,..., PhD	Pre-study survey
Job role	Open-ended	What is your work role?	Pre-study survey
Color-emotion associations	Open-ended color-emotion and emotion-color association tests.	1. You will now see swatches of different colors. For each color, please indicate what emotion(s) this color evokes, or how it makes you feel. 2. If you had to express your emotion via color, what color would you choose to indicate [stress]?	Pre-study survey
Affect: valence	Two-dimensional Circumplex model of affect (Russel 1980)	Score on a horizontal axis ranging from -100 (negative affect) to +100 (positive affect).	Daily experience sampling

Affect: arousal	Two-dimensional Circumplex model of affect (Russel 1980)	Score on a vertical axis ranging from -100 (low arousal) to +100 (high arousal).	Daily experience sampling
Engagement	Likert-type scale (Schaufeli, Salanova, González-Romá, & Bakker, 2002) anchored at 0=not at all to 5=extremely.	In the task you were just doing, how engaged were you?	Daily experience sampling
Challenge	Likert-type scale (Schaufeli et al., 2002) anchored at 0=not at all to 5=extremely.	In the task you were just doing, how challenged were you?	Daily experience sampling
Affect	PANAS, a 20-item inventory of positive and negative affect (Watson, Clark, & Tellegen, 1988); responses are on a 1=not at all to 5 = extremely scale	Indicate to what extent you feel this way right now, that is, in the present moment: [interested, excited, distressed, upset...].	End-of-day daily survey
Work day description	Likert-type scale anchored at 1=strongly disagree to 5= strongly agree.	I felt that I was under work pressure today. I felt this was a typical workday for me. I feel that something significant influenced my behavior at work today.	End-of-day daily survey
Work day: stress control	Likert-type scale anchored at 1=strongly disagree to 5= strongly agree.	I felt in control of my stress.	End-of-day daily survey
Work day: productivity	Likert-type scale anchored at 1=strongly disagree to 5= strongly agree.	I feel that I was productive in my work today. I felt in control of my productivity.	End-of-day daily survey
Job satisfaction	Index (Brayfield & Rothe, 1951); five-items anchored at 1=strongly disagree and 7=strongly agree	I feel fairly satisfied with my present job.	Pre- and post-study surveys
Crystal evaluation	12 items assessing various aspects of the Crystal use on a Likert-type scale anchored at 1=strongly disagree to 5= strongly agree	My Crystal was easy to read. My Crystal was very unobtrusive. My Crystal helped me to control my stress level.	End-of-day daily survey (week 2 only)
Crystal evaluation	Open-ended	Overall, how did you feel about using Crystal today? I see myself as someone who is [talkative, can be moody, is unselfish with others...]	End-of-day daily survey (week 2) Post-study survey
Personality traits: extraversion, openness, conscientiousness, agreeableness, neuroticism	Big 5 Personality inventory (John & Srivastava, 1999), 44-item instrument with each item measured on a 1=strongly disagree to 5= strongly agree	Likert-type scale	
Self-monitoring	Self-Monitoring Index (Snyder, 1974), 25 “true” or “false” items.	My behavior is usually an expression of my true inner feelings, attitudes, and beliefs.	Post-study survey
Emotional intelligence	EI Index (Schutte et al., 1998); 33 items anchored at 1=strongly disagree and 7=strongly agree	I have control over my emotions. I easily recognize my emotions as I experience them.	Pre-study survey

RESULTS

Affective states: physiological data

During the experimental week (i.e., week 2 of the study), participants collected 133.37 minutes of usable physiological data per day ($SD=62.49$), totaling over 115 hours of data from 47 person-days. Figure 5 shows that overall, participants were mostly in the neutral mode ($M=43.20$ minutes a day, $SD=30.42$), negative affect-low arousal mode (e.g., bored, fatigued, sad; $M=42.96$, $SD=25.54$), or negative affect-high arousal mode (e.g., stressed, frustrated, angry; $M=33.87$, $SD=18.19$). The Crystal recorded an average of only 7.06 minutes ($SD=9.48$) of positive affect-high arousal state (happy), and 6.28 minutes ($SD=8.61$) of positive affect-low arousal (e.g., calm, pleased) a day. Thus, according to the physiological data, approximately 90% of the time users were either in the neutral or in the negative affect modes.

These findings indicate that highly (and even moderately) positive emotions are not typical in the workplace. This is consistent with our informal observation – in general, knowledge and administrative workers are either in a relatively balanced, unemotional state if they are having a typical day and all goes well, or experience one of the common negative workplace emotions - stress, fatigue, or boredom. Highly positive events, such as research breakthroughs, paper acceptances, awards and promotions, are presumed to be less frequent. As we learnt from the interviews and informal verbal probes throughout the study, most of highly positive emotions were evoked by positive interactions with participants' colleagues.




Figure 5. The average daily distribution of the Crystal colors.

Affective states: self-reports

Recall that in parallel with physiological data collection, all participants were receiving experience sampling probes that asked them to self-report their arousal and valence. During the experimental phase of the study, our software collected 218 usable experience sampling probes, or an average of 4.34 probes per person per day (approximately one probe every half an hour). Since all probes were time-stamped, we can compare their results with the subset of physiological data collected the minute of the probe deployment. As shown on the plot (Figure 6), most of affect self-reports are concentrated in the positive valence-high arousal quadrant. Only 12% of all self-reports had a negative valence, and only 20% had a negative arousal value. To formally test a positivity bias, we conducted a univariate repeated-measures ANOVA test with self-reported and machine-predicted values of affect as within-subject variables. We found that predicted arousal values ($M=-12.03$, $SD=74.45$) were significantly lower than the self-reported ones ($M=24.33$, $SD=33.58$, $F(1,220)=47.31$, $p<.001$), and predicted valence values ($M=-46.10$, $SD=47.84$) were lower than self-reported valence values ($M=23.69$, $SD=25.73$, $F(1,220)=332.78$, $p<.001$). Although not the primary focus of the present study, this comparison was an interesting exercise with methodological implications. One possible explanation of the discrepancy between self-reported and physiological measures of affect is response bias, where participants' consciously or subconsciously rate their mood as more energetic and positive than it really is. Other studies have also found a bias toward positivity in the self-reported affect ratings (e.g., Mark et al. 2014), which may indicate that the phenomenon is systematic and not unique to our study.




Figure 5. Occurrence of all the predicted (blue dots) and self-reported (red dots) affective states detected during the study.

How accurate is the Crystal? Neither automatic affect assessments nor self-reports can be accepted as a ground truth, therefore we cannot compare data from these two sources to establish the Crystal's accuracy. However, one would hope that the two measures correlate. To test how physiological readings correspond to self-reports of valence and arousal, we ran the Pearson's correlation test for both of these variables. We found significant but small relationships: $r_{\text{arousal}}=.16$, $p<.05$; $r_{\text{valence}}=-.15$, $p<.05$. Inspection of individual results revealed that the relationship was significant for some participants, but not for all, and in some cases the relationship was negative.

Predicting challenge and engagement

Pop-ups that measured users' affective state also included two questions about how engaged and challenged participants were at the moment of measurement. Overall, our participants were moderately engaged ($M=3.18$, $SD=1.44$ on a 0-5 point scale) and challenged ($M=2.33$, $SD=1.55$) throughout the study. A series of Pearson's correlation tests showed that the predicted arousal was not related to how challenged or engaged the user was (at least at the level of arousal detectable by the Crystal; Table 2). Predicted valence was not related to the degree of engagement, but was negatively correlated with the challenge score – demanding tasks seemed to ruin people's mood. These results are consistent with those previously reported by (Mark, Iqbal, Czerwinski, Johns, 2014, April).

Self-reports, once again, showed different results. First, arousal was strongly correlated with both challenge and engagement. It is possible that the users simply made some inferences about their state, and concluded (if only implicitly) that their high arousal must indicate a high degree of engagement and involvement with the task. Alternatively, they could conclude that their strong engagement with a challenging task must have made them aroused.) There was also a small but significant relationship between self-reported valence and the degree of challenge. One possible explanation is a simple demand effect. Being challenged at work is generally associated with mental stimulation, competence, and determination, so informal rules of professional conduct dictate that challenging situations should be accepted favorably.

Table 2.

Relationship between challenge, engagement and two dimensions of affect.

	Challenge	Engagement
Predicted valence	-.25**	-.12
Predicted arousal	-.04	-.03
Self-reported valence	.15*	.05

Self-report arousal	.39**	.45**
---------------------	-------	-------

*p<.05 **p<.01

Crystal evaluation

What did our participants think about the BioCrystal? In order to test whether users' opinions changed over time (i.e., from day one to day five of week 2), we conducted a one-way repeated-measures ANOVA on subjects' responses. We did not find any effects of practice or habituation, meaning that participants' ratings of the BioCrystal remained stable throughout the study. For further analysis, we averaged responses across all days (Table 3). Overall, participants enjoyed using the BioCrystal. They felt the device was quite unobtrusive yet easy to read – both critical qualities for an ambient display. Users rated the Crystal's accuracy as average, and reported that it was moderately useful for controlling their stress levels.

Table 3.
Users' evaluation of the Crystal.

	M (SD)
My Crystal was too disruptive.	1.69 (0.39)*
My Crystal helped me to control my stress level.	2.82 (0.87)
I was not paying attention to my Crystal.	2.48 (0.94)
My Crystal showed exactly how I felt.	3.20 (0.69)
My Crystal was very unobtrusive.	3.32 (0.79)
I found my Crystal to be useless.	2.40 (0.92)
I was monitoring my Crystal throughout the day.	2.73 (0.84)
My Crystal signal seemed to be off mark.	3.37 (0.93)
Overall, I enjoyed using Crystal.	3.40 (1.07)
My Crystal was easy to read.	4.20 (0.44)
I was paying attention to other people's crystals.	3.20 (1.13)
Other people's crystals helped me understand their emotions.	2.40 (0.92)

*Responses are on a 1 (strongly disagree) to 5 (strongly agree) scale

To capture any indirect and subtle effects of the BioCrystal, we complemented users' ratings of the device with their self-reports of mood, stress level, job performance and job satisfaction. These variables were measured at the end of each day of the study, both in week 1 (control, no BioCrystal) and week 2 (experimental). First, we ran a series of one-way repeated-measures ANOVAs to test how users' responses changed from the beginning to the end of each week. We did not find any significant changes on any of the measures, and averaged participants' daily responses to compute one weekly score on each variable for each user. We then used a one-way repeated-measures ANOVA to compare users' responses in the control and experimental phases of the study.

The set of control questions (e.g., "This was a typical day for me", "I was under work pressure") confirmed that users' work conditions during both phases of the study were similar. Did the Crystal make any difference in participants' well-being and work outcomes? Having the Crystal did not seem to increase participants' self-reported productivity or overall job satisfaction, and did not change their self-reported mood. However, when using the Crystal during the second week of the study, participants reported having greater control of their stress ($M=3.61$, $SD=0.68$ on a 5-point scale) than on week 1 ($M=3.33$, $SD=0.67$, $F(1, 9)=4.41$, $p=0.06$) – arguably, the most important benefit of a mood-awareness device. Again, note that this finding is different from the one we obtained in response to the direct question about the stress-management properties of the BioCrystal (see Table 3). When we explicitly asked participants whether the Crystal helped them to control their stress level, they neither agreed nor disagreed ($M=2.82$, $SD=0.87$ on a 5-point scale). However, all participants reported feeling more control over their stress. Because in all other respects the two study weeks were non-different, it is reasonable to suggest that the effect is due to the use of the device. In other words, the Crystal did have a stress-management property, but our participants failed to recognize or acknowledge it.

It is also possible that some users simply did not have enough time and opportunity to experience and evaluate stress-control benefits of the Crystal. A Pearson's correlation test revealed that the longer the participant used the

Crystal, the higher she rated its usefulness for managing stress ($r=.77$, $p<.001$). Of course, one interpretation is obvious: those for whom the Crystal was useful, tended to use it more. However, the reverse is also possible: in order for the Crystal to prove its benefits, the user has to use the device for a sufficient amount of time. This is particularly true for participants low on neuroticism and high on self-monitoring. As we report next, these personality traits were related to the overall usage of the device.

The effects of dispositional variables and work conditions

We ran a series of Pearson's correlation tests to explore how the Crystal usage is related to work circumstances and participant's personality. Emotional intelligence was not related to any aspect of the Crystal use. Contrary to our expectations, we also failed to find any significant influence of the key personality traits (i.e., Big Five). The only exception was a strong positive relationship between neuroticism and perceived accuracy of the BioCrystal as measured by users' self-report: $r=.68$, $p<.05$. One possible explanation is that individuals high on neuroticism exhibit a wider range of affective states, which are easier for the sensors to detect.

Self-monitoring emerged as a significant predictor of the daily duration of the Crystal use (recall that participants could start and finish sessions at any time). Individual self-monitoring scores were negatively correlated with the average amount of time (in minutes) users spent with their Crystal: $r=-.69$, $p<.05$. This pattern is probably explained by the communication function of the Crystal. High self-monitors are, by definition, concerned with their self-presentation and strive to ensure appropriate public appearance. The Crystal's display of emotion – whether accurate or not – is based on physiological data and, therefore, is difficult to manage unless the user engages in deep acting. Because the BioCrystal can interfere with impression-management efforts and reveal an emotion that is inappropriate in the workplace (e.g., boredom), high self-monitors may be less willing to use the Crystal, or any other affect-signaling device with a visible display.

We also found that work conditions significantly influenced how participants used the BioCrystal. First, busy users did not pay attention to the device. The more deadlines the user had, and the less typical her day was, the less attention she paid to the Crystal ($r=-.69$, $p<.001$ and $r=.64$, $p<.05$, respectively). Being in a positive mood also made people less interested in the device ($r=-.64$, $p<.05$). Thus, the users most likely to monitor the Crystal were those having a typical workday, in a negative mood but not too stressed.

User insights

On the last day of the study all users were interviewed about their experience with the Crystal. We obtained several interesting insights about the benefits and potential applications of this affect-measuring technology.

Identifying stress-triggers. Several users noted that the Crystal not only made them more aware of their affective state, but also helped them to identify triggers of that state:

[The Crystal] allowed me to reflect on my mood, my state. It made me think about how I feel. For example, if [my Crystal] glows red, I stop and think: "Why? Am I stressed? What am I doing that is making me stressed?"

Certain things triggered different colors; [the Crystal] made me realize what triggers what.

The detrimental impact of stress on individuals and organizations is widely acknowledged, but very little attention is paid to the sources of stress and ways to identify it. Devices like Crystal could be effective in helping users to recognize when they are feeling stressed out and to develop positive personal coping strategies.

The ability of Crystal to identify stress-triggers has other practical ramifications that go well beyond mood-management. Specifically, being able to identify actions and objects that change the mood can be instrumental in campaigns aimed at breaking various health and consumer habits. Habits are known to be highly stable automatic behaviors that are triggered by cues in the environment, and revealing links between such environmental cues and user's actions and emotions is critical for behavioral change (Neal & Wood, 2007).

Facilitating mood management. Participants reported that the Crystal not only alerted them to their increased stress levels, but also helped them in their efforts to manage that condition:

I also found myself thinking about how I could move from one color to another. For example, what would I need to do [for my Crystal] to become green? Would I need to stand up and walk?

I was trying to change color from red, take a minute, and take a deep breath. Sometimes it worked.

Thus, by providing biofeedback in real time, the Crystal not only warns its user about her state, but also tracks what happens when the user is trying to change that emotional state. Arguably, this use of the Crystal can prove to be one of its most useful applications. Users interested in managing their stress have a variety of well-known options available to them – e.g., yoga, meditation, walks, just to name a few of the most popular ones. However, there is no technology that would help the user measure success of the mood-management methods and guide her in those efforts. And, we believe that we could see real improvements to this effect if we study this technology for a longer period of time.

Improving communication in the workplace.

Emotional expressions are critical for interpersonal communication. However, their effects depend on the observer's ability to read emotions (Van Kleef et al., 2012). In contrast to some other mood-trackers that display information only to its user (e.g., the commercial *MoodSense* phone app, Softonic), the Crystal light could be made visible to others by setting a transparent (vs. opaque) shell. By making emotional states more obvious, did Crystals facilitate interpersonal communication in the workplace? Our participants reported that this feature of the Crystal helped, or could help, their communication with colleagues or family:

I'm an introvert, I'm shy. I would be in white, and then someone walks in, and it turns red. It was ... a good way to signal that I may not be [available]. It was like a shield.

I would put [the Crystal] outside for coworkers to know what mood I am in, and what to expect from me. I would also want to see others' Crystals so that I could always know what is going on... the affective climate. Then I would know if I should start a meeting with a group hug, or if I could jump straight to the difficult topic, or maybe take somebody for coffee. So that I could tailor communication...

I was looking at other people's Crystals... We are administrative support team, and it is good to know how other people are doing at the moment, whether they are overwhelmed, and maybe need help. If I knew somebody was stressed, I would be gentler, would change how I talk to them (if at all). It would be also good to signal availability. And if all the team is "red", then maybe management needs to do something for the team, something on the group level. Maybe some relaxing activity.

I was always paying attention to other people's crystals, was more aware about what influence I had on other people... I did not want to disturb them.

Interestingly, participants' responses revealed that in the area of interpersonal communication, they treated the Crystal's signal similar to the natural display of emotions – they were mindful of what their Crystals revealed about them, and engaged in strategic management of that emotional display. Consistent with informal rules of professional conduct (Grandey, 2000; Kramer & Hess, 2002), one participant preferred to always keep her Crystal in the private mode, and three other participants switched between the private and the public modes:

At first I kept [my Crystal] [in opaque mode], but then I decided it would be best if people knew my state. But those were only people from my group. They are very friendly, and we were laughing about it, so I didn't mind displaying [my Crystal to] my group. But if talking to a manager, I would put it [in the private mode].

I always kept my own Crystal [in opaque shell]. I am an admin, I don't want people to see I am in bad mood, or have no energy, or stressed. I don't want to scare people off, or send them a wrong signal.

It would be good to have two separate [Crystals] – one for private use in the office, and one for interacting [with coworkers]. Something wearable... I would use the one in the office as a self-reflection [device] for stress-coping. It would really help me to be more mindful of what gives me stress. And I would use the wearable one as an interaction tool, to signal if I am available to chat, work etc.

CONCLUSION

We have presented the Crystal – a biofeedback device that measures users' affective state in real-time and communicates it via an ambient display. We made an effort to use real-world context and were able to achieve high ecological validity. This is important, as real-world settings could render the use of physiological signals problematic. Our findings from a two-week multi-method field study confirm that the Crystal helps to control stress, and can be successfully used to increase users' awareness of their affective state, which could lead to behavioral change. The Crystal also has the benefits of a mood-sharing device, and can be used to facilitate communication in the workplace and other settings.

Our study provides several main contributions to the research on ambient and affect-sensing devices. First, we demonstrate an innovative iterative research approach to the design of personal ambient affective display showcasing the relationship between design inspiration and empirical research.

Second, we provide evidence for the feasibility and utility of an ambient display that reflects the affective state of its owner. While many ambient displays have been explored over the last decade, most of the reports provide very little, if any, formal evaluation of the device usability. We conducted a systematic testing of the Crystal, and demonstrated that an ambient affect-monitor with a visible display has several benefits for the users. Arguably the most important benefits are affect-awareness and stress-management functions of the device. Although the Crystal did not influence the actual level of users' stress – that would be a pretty ambitious goal for any device, especially given a short time of the intervention – by the end of the study user's had greater awareness of their affective states, triggers of these affective states, and felt in greater control of their stress. Participants also noted the Crystal's ability to improve interpersonal communication by signaling others when it might be bad time to interrupt them.

Third, we explored the role of dispositional and contextual factors that can potentially influence usage of the affect-monitoring technology. In our study, neither emotional intelligence, nor Big Five traits predicted how participants evaluated device. However, we found that high self-monitors used the Crystal less than low self-monitors. Many of our users reported that visibility of the Crystal display made them mindful about what message they were conveying to the colleagues. Because self-monitors are particularly concerned with their expressive self-presentation, such a public display of emotion is incongruent with their efforts for expressive control. Therefore, high self-monitors used the device less. The use of the Crystal was also influenced by the situational factors. We found that individuals who were in a positive mood were less likely to monitor their Crystals than those in a negative mood. Busy and overly stressed participants also ignored the device.

LIMITATIONS

Our study design had several limitations which left some important questions unanswered. First, due to the sampling method, our pool of subjects was imbalanced in terms of gender. Research consistently finds significant differences in males' and females' affect, reactions to stress, and stress-management strategies. For example, women engage in emotional labor – i.e., strategic and effortful management of emotional display – more often and more successfully than men (Grandey, 2000). Women are also more likely than men to experience the negative effects of stress, and employ different coping mechanisms than men (Bickford, 2005). It has been found that, in general, women tend to use more social/emotional strategies to cope with stress, whereas men are more likely to use behavioral/mental or drug/alcohol disengagement. Thus, the Crystal may have different utility for men, they can evaluate its properties differently, and use it for different purposes. Researchers need to take caution in generalizing our findings to male users.

All our participants were also highly skilled knowledge workers with at least some college education. Knowledge workers have been successfully used as participants in other affect-tracking studies (e.g., Mark, Iqbal, Czerwinski, &

Commented [YL4]: I think this is an important part of the paper. Do we want to elaborate on it? MPC: I think its ok as is.

Commented [YL5]: I did not discuss two other findings: discrepancy between physiological and self-reported data, and correlation between engagement/challenge and affect. I feel like those findings are secondary and would distract from the main points of the work. What do you think? MPC: works for me.

Johns, 2014). However, we believe that work conditions and task characteristics of our participants (i.e., good familiarity with information technology, dealing with multiple tasks and deadlines at work) are only similar to work conditions of other information workers. We therefore can generalize our results to some, but not all, types of users. We suggest future studies to have a more inclusive, diverse sample of the participants.

Finally, we must acknowledge methodological challenges we faced in conducting this study. Testing ambient displays is always complicated by the conflict between the relatively intrusive evaluation and unobtrusive nature of the device that is intended to function at the periphery of the user's attention. Probing from the researchers constantly reminds the users about the device being tested, possibly increasing user's attention to the device and even prompting changes in emotional states. We tried to minimize intrusiveness of our measures by administering most of the questions at the end of the day or even at the end of the project. However, we do not know whether these efforts were sufficient. There is a general lack of research discussing the methodological difficulties of evaluating ambient displays. We hope that future studies will focus on these issues and will provide some viable solutions.

SUGGESTIONS FOR FUTURE RESEARCH

Based on results of our usability study, we identified two promising directions for future research on ambient affect-sensing displays. First, our participants expressed much interest in the possibility of using a wearable version of the Crystal (e.g., a bracelet, or a ring). Users mentioned that they would use such a device for private stress-management purposes, whereas desktop Crystal would be used as a tool for interpersonal communication. We agree that portability of the affect monitor would greatly increase its utility, and would enable its use in contexts other than seated workplace: homes, hospitals, children and adult day care centers, on the road by commuting drivers, etc.

Another interesting direction for the future work to explore is the type of feedback provided to the user. In the present study, the Crystal functioned as a biofeedback device that displayed actual affective state of the user. In the future, it would be useful to test the possibility of using such a monitor to induce certain affective states by providing modified feedback. For example, a stressed user would see green ("calm") color instead of red, a tired or a sleepy user would see bright blue or yellow, etc. A similar idea was executed and described by Stefani and colleagues (Stefani et al., 2011), but researchers did not conduct a formal usability test, and there is still lack of evidence about the effect such an intervention would have.

References

AlZoubi, O., D'Mello, S. K., & Calvo, R. A. (2012). Detecting naturalistic expressions of nonbasic affect using physiological signals. *Affective Computing, IEEE Transactions*, 3(3), 298-310.

Baas, M., De Dreu, C. K., & Nijstad, B. A. (2008). A meta-analysis of 25 years of mood-creativity research: Hedonic tone, activation, or regulatory focus? *Psychological Bulletin*, 134(6), 779.

Bickford, M. (2005). Stress in the workplace: A General overview of the causes, the effects, and the solutions. *Canadian Mental Health Association Newfoundland and Labrador Division*, 1-3.

Bono, J. E., & Vey, M. A. (2007). Personality and emotional performance: extraversion, neuroticism, and self-monitoring. *Journal of Occupational Health Psychology*, 12(2), 177.

Boucsein, W. (2012). *Electrodermal activity*. Springer.

Brayfield, A. H., & Rothe, H. F. (1951). An index of job satisfaction. *Journal of Applied Psychology*, 35(5), 307.

Brief, A., & Weiss, H. (2002). Organizational behavior: Affect in the workplace. *Annual Review of Psychology*, 53, 279-307.

Cacioppo, J. T., & Tassinary, L. G. (1990). Inferring psychological significance from physiological signals. *American Psychologist*, 45(1), 16.

Calvo, R. A., & D'Mello, S. (2010). Affect detection: An interdisciplinary review of models, methods, and their applications. *Affective Computing, IEEE Transactions*, 1(1), 18-37.

Chang, A., Resner, B., Koerner, B., Wang, X., & Ishii, H. (2001). LumiTouch: an emotional communication device. *CHI'01 extended abstracts on Human factors in computing systems* (pp. 313-314). ACM.

Church, K., Hoggan, E., & Oliver, N. (2010). A study of mobile mood awareness and communication through MobiMood. In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries* (pp. 128-137).

Cropanzano, R., James, K., & Konovsky, M. A. (1993). Dispositional affectivity as a predictor of work attitudes and job performance. *Journal of Organizational Behavior*, 14(6), 595-606.

D'Mello, S. K., Craig, S. D., Witherspoon, A., McDaniel, B., & Graesser, A. (2008). Automatic detection of learner's affect from conversational cues. *User modeling and user-adapted interaction*, 18(1-2), 45-80.

Ekman, P., & Friesen, W. V. (2003). *Unmasking the face: A guide to recognizing emotions from facial clues*. Ishk.

Forgas, J. P. (1995). Mood and judgment: the affect infusion model (AIM). *Psychological Bulletin*, 117(1), 39.

Frijda, N. H., & Mesquita, B. (1994). The social roles and functions of emotions. In S. Kitayama & H. S. Markus (Eds.), *Emotion and culture: Empirical studies of mutual influence* (pp. 51-87). Washington, DC: American Psychological Association.

Gangestad, S. W., & Snyder, M. (2000). Self-monitoring: appraisal and reappraisal. *Psychological Bulletin*, 126(4), 530.

Gay, G., Pollak, J. P., Adams, P., & Leonard, J. P. (2011). Pilot study of Aurora, a social, mobile-phone-based emotion sharing and recording system. *Journal of Diabetes Science and Technology*, 5(2), 325-332.

Geven, A., Tscheiligi, M., & Noldus, L. (2009). Measuring Mobile Emotions: Measuring the Impossible?. In *Proceedings of the 11th International Conference on Human-Computer Interaction with Mobile Devices and Services* (p. 109). ACM.

Golder, S. A., & Macy, M. W. (2011). Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science*, 333(6051), 1878-1881.

Grandey, A. (2000). Emotion regulation in the workplace: A New way to conceptualize emotional labor. *Journal of Occupational Health Psychology*, 5 (1), 95-110.

Hazlewood, W. R., Stolterman, E., & Connelly, K. (2011). Issues in evaluating ambient displays in the wild: two case studies. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 877-886). ACM.

Hektner, J. M., Schmidt, J. A., & Csikszentmihalyi, M. (Eds.). (2007). *Experience sampling method: Measuring the quality of everyday life*. Sage.

Herzberg, F. (1993). *The Motivation to work*. Transaction publishers.

Hochschild, A. R. (1979). Emotion work, feeling rules, and social structure. *American Journal of Sociology*, 551-575.

John, O. P., & Srivastava, S. (1999). The Big-Five trait taxonomy: History, measurement, and theoretical perspectives. In L. A. Pervin & O. P. John (Eds.), *Handbook of personality: Theory and research* (Vol. 2, pp. 102-138). New York: Guilford Press.

Kapoor, A., & Picard, R. W. (2005). Multimodal affect recognition in learning environments. In *Proceedings of the 13th annual ACM international conference on Multimedia* (pp. 677-682). ACM.

Kramer, M.W., & Hess, J.A. (2002). Communication rules for the display of emotions in organizational settings. *Management Communication Quarterly*, 16(1), 66-80.

Liu, K. K. L. (2004). *A personal, mobile system for understanding stress and interruptions* (Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA).

Loewenstein, G., & Lerner, J. S. (2003). The role of affect in decision making. *Handbook of affective science*, 619 (642), 3.

MacLean, D., Roseway, A., & Czerwinski, M. (2013). MoodWings: a wearable biofeedback device for real-time stress intervention. In *Proceedings of the 6th International Conference on Pervasive Technologies Related to Assistive Environments* (p. 66). ACM.

Madden, T., Hewett, K., & Roth, M. (2000). Managing images in different cultures: A Cross-national study of color meanings and preferences, *Journal of International Marketing*, 8 (4), 90-107.

Mark, G., Iqbal, S., Czerwinski, M., & Johns, P. (2014a). Capturing the mood: Facebook and face-to-face encounters in the workplace. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing* (pp. 1082-1094). ACM.

Mark, G., Iqbal, S., Czerwinski, M., & Johns, P. (2014b). Bored Mondays and focused afternoons: the rhythm of attention and online activity in the workplace. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems* (pp. 3025-3034). ACM.

McDuff, D., Karlson, A., Kapoor, A., Roseway, A., & Czerwinski, M. (2012). AffectAura: an intelligent system for emotional memory. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 849-858). ACM.

Neal, D. T., & Wood, W. (2007). Automaticity *in situ*: Direct context cuing of habits in daily life. In J. A. Bargh, P. Gollwitzer, & E. Morsella (Eds.), *Psychology of action (Vol. 2): Mechanisms of human action*. London: Oxford University Press.

Olivera, F., Rivas, A., & Iturriaga, F. (2013). Subtle Interaction for a Non Intrusive Communication. In *Ubiquitous Computing and Ambient Intelligence. Context-Awareness and Context-Driven Interaction* (pp. 215-222). Springer International Publishing.

Picard, R. W., Vyzas, E., & Healey, J. (2001). Toward machine emotional intelligence: Analysis of affective physiological state. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 23(10), 1175-1191.

Pillutla, M. M., & Murnighan, J. K. (1996). Unfairness, anger, and spite: Emotional rejections of ultimatum offers. *Organizational Behavior and Human Decision Processes*, 68(3), 208-224.

Rea, D. J., Young, J. E., & Irani, P. (2012). The Roomba mood ring: an ambient-display robot. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction* (pp. 217-218). ACM.

Russell, J. (1980). A Circumplex model of affect. *Journal of Personality and Social Psychology*, 39, 1161-1178.

Salovey, P., & Mayer, J. (1989). Emotional intelligence. *Imagination, Cognition, and Personality*, 9(3), 185-211.

Schaufeli, W. B., Salanova, M., González-Romá, V., & Bakker, A. B. (2002). The measurement of engagement and burnout: A two sample confirmatory factor analytic approach. *Journal of Happiness Studies*, 3(1), 71-92.

Schutte, N. S., Malouff, J. M., Hall, L. E., Haggerty, D. J., Cooper, J. T., Golden, C. J., & Dornheim, L. (1998). Development and validation of a measure of emotional intelligence. *Personality and Individual Differences*, 25(2), 167-177.

Snyder, M. (1974). Self-monitoring of expressive behavior. *Journal of Personality and Social Psychology*, 30(4), 526.

Softonic MoodSense mobile phone application <http://bit.ly/9uKEVZ>, last retrieved on 5/28/2014.

Stefani, O., Mahale, M., Pross, A., & Bues, M. (2011). SmartHeliosity: emotional ergonomics through colored light. In *Ergonomics and Health Aspects of Work with Computers* (pp. 226-235). Springer Berlin Heidelberg.

Van Kleef, G., Homan, A., & Cheshin, A. (2012). Emotional influence at work: Take it EASI. *Organizational Psychology Review*; doi: 10.1177/2041386612454911

Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: the PANAS scales. *Journal of Personality and Social Psychology*, 54(6), 1063.

Watson, D., & Clark, A.L. (1992). On traits and temperament: General and specific factors of emotional experience and their relation to the five-factor model. *Journal of Personality*, 60(2), 441-476.