Health Chair: Implicitly Sensing Heart and Respiratory Rate

Erin Griffiths
University of Virginia
Charlottesville, VA, USA
erg3wb@virginia.edu

T. Scott Saponas
Microsoft Research
Redmond, WA, USA
ssaponas@microsoft.com

A.J. Bernheim Brush
Microsoft Research
Redmond, WA, USA
ajbrush@microsoft.com

ABSTRACT

People interact with chairs frequently, making them a potential location to perform implicit health sensing that requires no additional effort by users. We surveyed 550 participants to understand how people sit in chairs and inform the design of a chair that detects heart and respiratory rate from the armrests and backrests of the chair respectively. In a laboratory study with 18 participants, we evaluated a range of common sitting positions to determine when heart rate and respiratory rate detection was possible (32% of the time for heart rate, 52% for respiratory rate) and evaluate the accuracy of the detected rate (83% for heart rate, 73% for respiratory rate). We discuss the challenges of moving this sensing to the wild by evaluating an in-situ study totaling 40 hours with 11 participants. We show that, as an implicit sensor, the chair can collect vital signs data from its occupant through natural interaction with the chair.

Author Keywords

Sensing chair; implicit sensor; vital signs sensing; resistive pressure sensor; EKG sensing

ACM Classification Keywords

J.3. Computer Applications: Life and Medical Sciences

INTRODUCTION

The most direct way to measure the general health of an individual is through their vital signs: including heart rate, respiratory rate, blood pressure, and oxygenation. However, most people have these vitals taken very rarely, perhaps at an annual doctor's visit or when an illness has already shown itself through other signs or symptoms. Though a motivated subset of the population may measure their vital signs during exercise, this is often restricted to periods of exercise only. Most of the population gathers only infrequent measurements of their vital signs.

While a recent explosion in wearable health devices [6, 16] enables more people to track their vital signs, there are a number of reasons people may choose not to wear such devices including comfort, appearance, and frequency of charging. Given the current costs and requirements of such devices, people adopting them are likely either highly motivated by a specific task, such as managing weight, or already aware of a chronic condition that requires monitoring. In contrast to these explicit health monitoring devices, we are interested in implicit health sensing. How can we sense vitals without any active participation from the individual being monitored?

Sensing vitals implicitly as people go about their daily lives requires the sensing to be embedded in objects that people interact with frequently. Research has shown that people currently spend most of their time in sedentary activities, sitting at work, sitting watching TV, sitting in cars, sitting at meetings, to eat, to socialize, etc. [8]. Many of these sedentary activities revolve around an object seen in nearly every home, building, and workspace: the chair.

Given their ubiquity, chairs are an intriguing possibility for implicit sensing. A sensing chair could tell what posture a person is holding, what their heart rate is, or how quickly they are breathing. It could monitor stressful reactions over the course of the day or record how the occupant's health changes over months or years. By combining ubiquitous sensing and the common, everyday chair, we envision a way to monitor a person's health between clinic visits and detect early signs or trends that can play a role in predicting future health.

To study the feasibility of implicit health sensing from a chair, we selected two vital signs to sense: respiratory rate and heart rate. These are commonly measured signs for many adverse conditions including anxiety, stress, sleep disorders, cardiovascular disease, and respiratory disease. While we are not the first to consider the possibility of a chair as a health sensor [e.g. 2, 18, 26], our unique contributions are:

- Qualitative data from 550 participants about their usage of chairs to inform the design of chair sensors
- Iterative design of an EKG armrest electrode pattern to increase implicit sensing opportunities
- Lab evaluation of respiratory and heart rate sensing accuracy across 7 common sitting positions with 18 people
- In-situ evaluation of implicit sensing using the Health Chair with 11 participants

Our findings suggest that the Health Chair can provide its occupants with information about their vital signs. Our experience moving the Health Chair out of the laboratory illuminated challenges that prevent the Health Chair from continuously sensing vitals when occupied. The Health Chair does, however, have the potential to be a ubiquitous implicit sensor for an individual's health, opportunistically measuring their vitals on a daily or multiple-times-per-day basis.

UbiComp '14, September 13 - 17, 2014, Seattle, WA, USA

ACM 978-1-4503-2968-2/14/09...$15.00.
http://dx.doi.org/10.1145/2632048.2632099

RELATED WORK

Using devices outside the clinic to measure the health and activities of people has been a growing area of interest in recent years [20]. These devices can roughly be separated into two main classes: wearable or explicit devices and environmental or implicit devices.

Explicit health devices are those worn by the individual being monitored. Many devices of this type are already commercially available and monitor a variety of health signs and activities, including devices that count the number of steps you take or monitor heart rate through a strap worn around your chest. Recent research has looked into making these devices smaller, cheaper, easier to wear, and easier to communicate with and collect data from, as well as how to collect a wider variety of health signs [e.g. 6, 16].

However, though these explicit devices measure many of the vital signs of the human body, they are not often used in everyday living. This is due to a number of problems with wearable devices. First, wearing devices can be cumbersome. If the device is too large, visible, or unfashionable individuals may be unwilling to wear the device. Second, monitored individuals may take the device off for activities such as showering, sleeping, dish washing, etc. and forget to put the device back on. Assuming the individual requires the device for immediate health purposes, such as a cardiac patient being monitored for signs of a heart attack, these first issues may be ignored or carefully worked around. However, this limits the monitored population to those with already diagnosed issues or identified high risk patients. Otherwise healthy individuals or those with unidentified health risks are almost never monitored by these systems.

Implicit sensing devices solve a number of the problems produced by explicit devices. As implicit devices are not worn, but instead embedded in the environment, there is no need for the individual to remember to don the device. Additionally, implicit sensing devices require no active participation from the monitored individual. Such devices include video monitoring to determine heart rate [17], weight sensors placed in beds [21], and sleep evaluation through heart rate and movement detection in beds [11]. A monitored individual can go about their daily life without interruption and each time they interact with these implicit sensing devices a snapshot of their health is taken.

The main drawback to implicit sensing then becomes the fact that it is not continuous. An individual must interact with or be in proximity to an implicit device for data on their health to be sensed. However, many of the devices endowed with sensing ability today are not in constant contact or sufficient proximity to a person for frequent measurement. Instead, implicit health sensing requires an object of great ubiquity and constant use as a sensing platform. In modern society, we can find a great candidate object below us during much of our waking hours: the common chair.

Recent research has examined the chair as an implicit sensor for several applications. Posture recognition through pressure sensors on the chair provides the basis for many of these applications [13, 12, 15, 23]. Applications that use posture recognition include: providing feedback to the user to correct bad posture [26], using positional information as a controller and user interface [22], and detecting the interest level and activities of the occupant [14, 10]. Additionally, the Smart Chair [4] uses pressure sensors mounted under the legs of a chair to extract information about the activities of the occupant, such as their postures and hand and head movements, to recognize activities such as typing on a computer and eating. These posture based approaches provide some indirect information about an individual's health, such as when their postures change frequently due to discomfort [2], but do not directly monitor the vital signs of the occupant.

Implicit sensors other than chairs that focus on vital sign detection include respiration and heart rate in beds [3, 11], blood pressure in toilet seats [24], and heart rate in car seats [7]. Ford has recently been developing heart rate sensors embedded in the backrest to measure heart rate through certain clothing materials. Toyota is developing a heart rate monitor built into a steering wheel [25], and Anttonen et al. [1] use an electromechanical film sensor (EMFi) covering a chair to detect heart rate. None of these sensors monitor respiration rate. Some implicit sensors monitor individuals while they are in chairs, but are not actually part of the chair, such as video monitoring to determine heart rate [17]. In this work we are looking to complement these approaches by providing a second sensing point when sensing capabilities overlap and sensing in situations these sensors cannot, such as when the occupant sits in the chair, but is facing away from the video monitor.

Perhaps most closely aligned with our approach, Postolache et al. [18] mounted an EMFi on the backrest and seat of a chair to monitor vital signs. Their results focus on evaluating a wavelet based signal processing algorithm to extract respiration and heart rate. Results presented from a lab study with 10 minutes of data collected for 10 people are encouraging, although details on sitting position and robustness of the approach across different positions is not provided so it is difficult to know how the approach would perform in real world conditions.

Our unique contributions that go beyond prior work include a chair usage study to inform our sensor design, lab evaluation across common sitting positions, and an in-situ evaluation of implicit sensing using a chair.

CHAIR USAGE STUDY

How successfully chairs can be used as implicit sensors depends on how frequently people are sitting, understanding how they sit to enable sensor placement, and the diversity of chairs people use throughout the day. To evaluate these factors we began our investigation with a usage study. Initially, we interviewed 50 office workers (M: 25 F: 25) with desk-based offices at our company. Before interviewing a participant we unobtrusively observed their current sitting position, including information such as how far forward they were sitting, if they were using their armrests, and how their legs were placed. Afterwards, if a participant chose to participate, this information was transferred to our participant pool and we asked them questions such as:

- Siting duration for their typical work day
- What features (e.g. armrests, backrests) they prefer on their chairs
- Typical posture they use in their chair

In addition to these interviews, we deployed a survey based on these questions on Amazon's Mechanical Turk (MTurk) to 500 participants (194 female) located in the U.S and paid each respondent U.S. $0.25. Due to the nature of the remote survey, Mturk participants were not observed in their natural pose before answering questions. However, the Mturk survey responses allow us to generalize our results beyond the initial office worker interviews. Of the Mturk participants, 59% spent less than 6 hours working on a computer, indicating that the Mturk respondents had more diverse jobs then than our internal population. In total, we collected data from 550 respondents (M: 331 F:219). The majority of participants were between the ages of 20 and 39 (75%), with 7% of participants under 20 and 17% above 40.

How much do we sit?

Prior work [8] tells us that in developed countries people spend up to 15 hours of their waking time in sedentary behavior. To better determine how much of this time was spent sitting, we asked all respondents to report the number of hours they spend sitting on a typical work or weekday. Participants were asked to think of all the times they sit during the day, including when sitting in cars, on buses, at mealtimes, at home, and at work. More than half of the participants (55%) reported that they typically spend more than 9 hours a day sitting. Of these, 20% reported they spend more than 14 hours a day sitting. Only 18% of respondents reported that they spend less than 6 hours sitting per day. These results indicate that, despite diverse jobs and lifestyles, people spend a large portion of their day sitting in chairs.

Where do we sit?

To better understand how many chairs would need to be instrumented to sense the majority of a person's waking day, we asked participants if they had a primary chair that they use during the day. 91% (498) of respondents reported they had a primary chair. Of these respondents, 61% reported that they typically sat in it for 6 or more hours a day. Surprisingly, 65% of respondents reported that this primary chair was located in their homes. These responses show that the majority of respondents interact primarily with a single chair, indicating that only one chair in their lives need be instrumented as an implicit sensor to achieve frequent sensing.

Although a person may have a primary chair, that chair may be used by multiple people. If multiple people use the same chair, implicit health sensing becomes more complicated because the chair must also identify the person sitting in it. We asked respondents whether or not they shared their primary chair with other people in order to determine how important it would be for the Health Chair to identify its occupant. Most of the respondents (69%) reported that they were the only ones to routinely use their main chairs, 18% report that only one other person did, and only 13% reported that 2 or more people did. Hence, for the majority of respondents, the chair would not need to determine their identity. Thus we do not focus on occupant identification in this paper.

How do people sit in chairs?

To understand how to sense vitals from someone in a chair we first need to better understand how people sit in chairs. For example, if an individual always leans forward when sitting in a chair and never uses the backrest, or does not have a backrest on their chair, then sensors designed to sense the person from the backrest would never see that individual. Sensing that requires skin contact, such as EDA or EKG, likely needs to be placed on armrests, since uncovered arms are more common in work and everyday settings than an uncovered back or upper legs.

To determine where sensing might be feasible on a chair we asked respondents how they commonly use the chair they sit in. The majority of respondents stated that they often use the backrests of chairs (67%) and only 6% reported that they never use backrests. Armrests were used, though to a lesser degree, with 38% reporting that they often used armrests and 42% reporting that they sometimes used armrests. Of the respondents that had a primary chair, 94% reported that their chairs had a backrest and 71% reported that they had armrests. Of the total respondents, 84% reported that they preferred armrests on their chairs and 87% reported that they preferred chairs with backrests. These numbers indicate that, for the majority of respondents, sensors in the backrest and armrests of their chairs would be in contact with the person's body for some of the periods during the day that they occupy their chair. Therefore we focus our sensing efforts on the armrests and backrest of the chair.

SENSING VITALS FROM A CHAIR

Informed by the chair usage study, we built our Health Chair prototype. We focused on sensing two vitals, heart rate and respiratory rate, as these are well known indicators for a variety of medical issues including cardiac and respiratory issues and other more subtle conditions such as stress and anxiety issues.

Our usage study showed that sensing from the armrests and backrest of chairs would provide periods of health sensing data throughout the day for most people. Therefore we chose these two locations from which to detect vital signs. Any sensing performed by an implicit sensor must be performed in such a way as to not impede the normal function of that object. In the case of the chair, this meant that sensors had to conform to the structure of the chair in question and produce no increase in discomfort for the occupant. Figure 1 shows the final prototype of the health chair, including the placement of the electrodes and pressure sensors. Our two embedded data collection devices are affixed out of sight to the bottom of the chair.

We now describe the hardware implementation and signal processing for heart rate and respiratory rate sensing.

Figure 1. We employ conductive fabric on the chair's armrests to sense heart rate and pressure sensors on the back of the chair for sensing respiratory rate (left). Two iterations of our armrests (right).

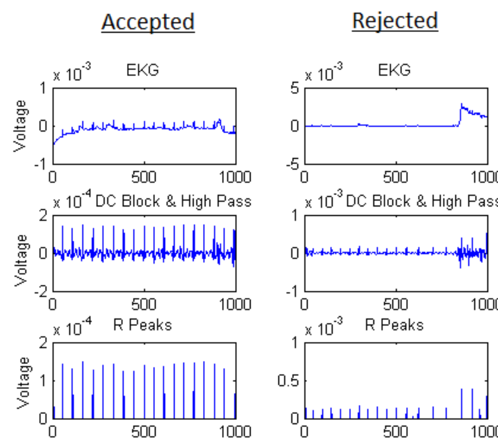

Figure 2. The rejected EKG signal (right) shows movement artifacts and variation in R peaks that is not seen in an accepted (left) signal.

Heart Rate from EKG

Clinicians commonly perform heart-rate detection using an electrocardiogram (EKG). An EKG measures the electrical activity of the heart, capturing the contraction of the heart muscle as blood is pumped throughout the body. Electrodes placed on the body capture these electrical signals through contact with skin. The Health Chair sensor detects this electrical activity through a three-electrode EKG, where two electrodes are placed on opposing sides of the heart to detect electrical activity and a third electrode drives the skin to a common voltage (referred to as the driven right leg or DRL) to keep the common-mode voltage of the two sensing electrodes in range of the analog-to-digital converter (ADC). While other research [7] is studying the accuracy of EKG recording without skin contact, our Health Chair requires skin contact.

Because EKG requires skin contact, electrodes were placed on both armrests of the chair, where an occupant is most likely to have exposed skin. We implemented the three electrode EKG used by the Health Chair using conductive fabrics as electrodes. The EKG electrodes are designed to cover the entire upper surface of the armrest. Each armrest is covered with two non-touching pieces of conductive material. Two iterations of the EKG electrodes were used in this study. The first used copper foil tape and silver ribstop of 0.25 Ohm/sq. The second used copper taffeta of 0.05 Ohm/sq and the silver ribstop. A combination of two materials was chosen for aesthetic reasons. A single material can be used for all electrodes provided electrodes are separated by an insulator. Both iterations of the armrests are shown in Figure 1.

The electrodes are designed with an interlocking zig-zag pattern. This pattern is intended to capture a person's arm or elbow with both pieces of fabric in any location where they come into contact with the armrest. Because the DRL electrode requires only one connection to the occupant, it is not necessary for the occupant to touch this contact with both arms (although it is fine if both arms touch DRL). Therefore, the DRL contact was placed as the electrode on the outside edge of the armrests. If an occupant places their arm on top of one armrest and leans into the other, touching the inside edge of the armrest only, this configuration will ensure that both opposing leads are connected to the occupant and able to collect EKG.

After our laboratory study, we learned that the individual stripes of conductive material were so thick that often a person would only be in contact with one of the active electrodes on an armrest. To address this issue, we made a second version of the armrest for the in-situ study with a tighter zig-zag pattern (see Figure 1).

We employ the Texas Instruments ADS1298 for EKG sensing. It is a low-power, 8-Channel, 24-Bit ADC meant for many types of biopotential measurements. The ADS1298 is from a family of biopotential ADCs suitable for EKG. We sample the electrodes at 100Hz and log the data to a nearby machine over Bluetooth.

Signal Processing

The signal processing component of the Health Chair extracts the rate of each vital sign being measured. Heart rate was calculated using a method called R-peak detection. The R wave represents the electrical stimulus in the heart as it passes through the main portion of the ventricular walls and is measured as the highest peak in a single beat shown in an EKG. R-peak detection detects this peak in each heart beat shown by an EKG signal. With this information the heart rate of the individual producing the EKG signal can be calculated. The method we used to detect R-peaks is based on the code published in [19].

To detect the R-peaks from the raw signal we begin by applying a first order IIR DC blocking filter with a cut-on frequency of 0.1 Hz. The signal is then divided into 10 second windows, with a one second step between windows. For each window, we apply a simple high pass by taking an FFT and then an IFFT of the frequency bins above 40 Hz. This sharpens the peaks in the signal that correspond to R-peaks (when there is good signal). From the high-passed signal, we can then find the local maximums in the signal which correspond to the R-peaks of the EKG signal.

Heart rate is calculated from the detected R-peaks by taking the average RR-interval (the time between two R-peaks) and inferring the heart rate of the individual in beats per minute. Once the inferred heart rate at a given second is found the rate is either accepted by the Health Chair as a rate corresponding to its occupant, or rejected due to inconsistencies, see Figure

2 for an example. Determining acceptable inferred rates is particularly important for implicit sensing because we expect the sensor to not always be in contact with the individual or to be detecting incorrect rates due to an occupant's movement, position, or other noise. An inferred heart rate can therefore be rejected as the occupant's heart rate for a number of reasons, most commonly because the individual is not in contact with the sensor. For the EKG sensor, lack of good contact also includes when only one of the occupant's arms is in contact with the armrests. A heart rate can also be rejected when there are enough inconsistencies in the signal to indicate that the detected rate may not be accurate. If the standard deviation of the RR-intervals in the 10 second window is greater than or equal to 0.1 seconds then the rate is rejected. Only accepted heart rates are attributed to the occupant.

Respiratory Movement

Respiratory rate is measured by observing the rise and fall of the human chest or abdomen over a period of time or the pressure exerted by these areas of the body expanding against adjacent surfaces. Because individuals breath differently, expanding through either their chest or their abdomen, no single location on the body exhibits respiratory expansion on all individuals. However, some portion of the human upper body always expands with respiratory movement making the backrest an ideal place to measure respiration using pressure sensors.

Respiratory movement was collected using force-sensing resistors attached to the backrest of the chair. A force-sensing resistor changes resistance when force is applied. The sensors used in the Health Chair are FSR 406 Square and FSR 408 Strip Force Sensing Resistors from Interlink Electronics. An Arduino attached to the bottom of the chair collected readings from each sensor. Data was logged onto a laptop connected to the Arduino over USB.

Most people expand or contract their upper body in one of two locations when they breath: their chest or their abdomen. The strip sensors are placed along the top of the chair to catch the expansion of the occupant's chest as breathing tilts there upper body back and forth. The square sensors on the lower back of the chair were placed to catch the expansion of the occupant's abdomen as breathing pushes the fleshy skin below their ribs out and in. It is important to note that an increase of pressure against the sensors does not indicate that the occupant has breathed in. Inhaling can lift parts of the occupant's body up and away from the sensors, particularly in the upper back, and exhaling can relax them back into the sensors. A total of six pressure sensors are used to detect respiratory rate as shown in Figure 1.

Signal Processing

Using a sliding 10 second window, with a one second step, we calculate respiratory rate from the pressure sensors using autocorrelation. Autocorrelation overlays the original signal on top of itself and shifts it across the time domain, measuring cross-correlation of the signal with itself. The first peak of similarity represents the lowest amount of lag time or movement before the signal matches itself. Movements of a person's chest as they breath in and out can often be seen as a sign-like wave in the pressure sensors, as they press into and pull away from the sensor. Thus, this first peak of similarity represents the frequency of the person's respiration.

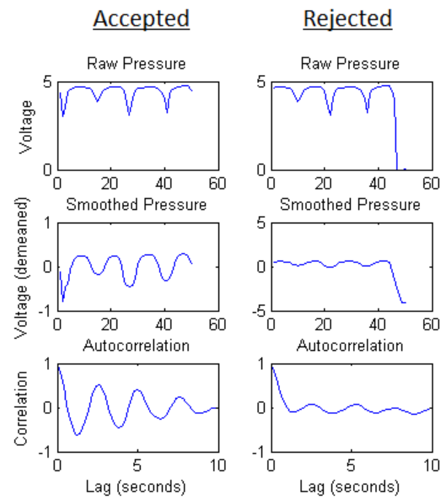

Figure 3. The accepted respiration signal(left) shows a much larger difference between the first valley and peak of the autocorrelation than the rejected signal (right).

As with heart rate, a inferred respiratory rate can be rejected for a number of reasons, including when the occupant is not in contact with the sensors. This is fairly common in the pressure sensors on the back of the chair, where a person might be in contact with some, but not all, of the sensors. A respiratory rate is rejected under two circumstances. First, if the magnitude of the frequency peak, minus the magnitude of its neighboring valleys is too small (less than 0.6), the rate is rejected. This means that any respiratory signal that does not show enough change between an out breath and an in breath is removed. Second, a rate is rejected when the peak-to-peak amplitude of the raw pressure signal is too small. This helps both to remove any calculated rates where the difference between the breath-in pressure and the breath-out pressure are not clear enough to indicate that the individual is causing them; instead, the sensor may be transducing noise sources such as slight deformation of the sensors in how they are attached to the chair. Figure 3 shows examples of accepted and rejected inferred rates.

Once a respiratory rate is inferred for each of the six pressure sensors, the final rates are fused to form the respiratory rate detected by the chair. Each respiratory rate that was accepted from each pressure sensor in a given 10 seconds window is averaged together to form the final inferred rate. The final respiratory rate is only accepted and attributed to the occupant if the standard deviation of these combined rates is less than or equal to 2 breaths per minute (found experimentally), ensuring that the pressure sensors closely agree on the final accepted respiratory rate.

LABORATORY EVALUATION

To assess how well the Health Chair could accurately detect heart rate and respiratory rate we first evaluated our chair in a laboratory study. Specifically, we looked at when the chair

Sitting Position	Arm Position	Back Position
1. Arms Flat	Arms Flat and parallel on the Armrests	Straight Back
2. Right Partial	Right Arm Resting Partially on Armrest, Left arm as in (1)	Straight Back
3. Arms Partial	Both arms resting partially on armrests	Straight Back
4. Right Elbow	Right elbow on armrest, Left arm as in (1)	Straight Back
5. Elbows	Both elbows on armrests	Straight Back
6. Slouched	Both arms as in (1)	Slouched, Lower back only touching
7. Leaning Back	Both arms as in (1)	Leaning back, Upper back only touching

Table 1. Participants varied both arm and back positions when using the chair in the laboratory study.

is able to detect these vital signs and how accurate is the detected rate?

Experimental Design

This study focused on assessing the ability of the Health Chair to accurately detect and calculate heart and respiratory rate under controlled conditions. We tested two aspects: 1) whether the Health Chair could detect a vital sign consistently and accurately when the vital is in a steady state, such as the heart rate while a person is at rest across several different sitting positions and 2) whether the Health Chair can detect when the vital sign is changing.

We recruited 18 participants (F:9, M:9) who sat in 7 different positions for 30 seconds each while silently watching a short video clip. These experiments tell us how accurately a rate at rest can be detected given how the occupant is sitting. The various positions were chosen from the observation of participants in our office building during the chair usage study. Participants were shown pictures of each position for reference, but were otherwise not directed into the position. In the first position, participants were asked to sit up straight with their back resting against the backrest and their forearms flat on the armrests. From this position participants were asked to vary their arm positions, as described in Table 1. Participants then returned to the first position and were asked to vary their back positions. To test how well the chair could detect when vital signs are changing we asked participants to bike on a stationary bike for a full minute before sitting in position 1 for a full minute. This exercise-sitting cycle was repeated four times.

We designed this study to evaluate the ability of the sensors to infer vitals when the occupant was in contact with the chair under controlled and ideal conditions. Thus, we asked participants to remove any jackets, long sleeves, or jewelry from their forearms that could interfere with the heart rate electrodes and we did not include sitting positions that would prevent contact with the sensors (e.g. not using the arm rests or sitting at the front edge of their chair and not touching the backrest).

Ground Truth

Ground truth for both heart rate and respiratory rate was collected using commercially available wearables. Heart rate and respiratory rate were detected using the Garmin Soft Strap Heart Rate Monitor and NeuLog Respiration Monitor Belt Sensor respectively.

The Garmin Soft Strap Heart Rate Monitor encircles an individual's chest and wirelessly transmits detected heart rates to a ANT+ device. The Garmin sensor requires good skin contact to detect heart rates and advises wearers to wet the strap for better conductivity before placing it on skin. If the strap does not have good conductivity it does not transmit a heart rate and because of this heart rate ground truth was lost for parts of the experiments for a number of the participants. The total amount of data lost across all participants and experiments was 51 minutes (36%). Only the parts of the experiments that had ground truth were evaluated. The Garmin strap has no specific sampling rate for detecting heart rates, whenever the monitor detects a heart rate it transmits it. The exact signal processing of the Garmin strap are proprietary. In order to match the heart rate detection of our chair, all heart rates detected by the Garmin strap in a single second are averaged together.

The NueLog Respiration Monitor Belt Sensor encircles an individual's upper abdomen and lower rib cage. An air bladder inside the belt is inflated to press against the chest of the individual and the air pressure within the bladder is recorded. The bladder is small, only 4-5 inches horizontally, and is positioned over the lower rib cage and upper abdomen of a participant's chest. The bladder itself does not press against the Health Chair's pressure sensors that touch a participant's back. When a person breathes their expanding chest is pressed against the bladder and the increase in pressure indicates an inhale. Before data is collected from each participant, the bladder is inflated until an inhale and exhale by the participant are seen as a sign wave in the pressure signal, then the bladder is stoppered to hold that pressure level. Since the NueLog senor will pick up all movement of the chest, participants were asked to refrain from talking and laughing so that the NeuLog only recorded the movements of their chest due to respiration. The resulting pressure signal was turned into respiratory rate by taking an autocorrelation of the signal and extracting the rate at the first peak of high correlation. As this is the same technique that is used to extract the respiratory rate from the Health Chair our comparison against respiratory ground truth is not measuring the error of the autocorrelation technique, but the similarity of the respiratory signal found by the Health Chair to that of the NeuLog Belt Sensor.

Results

The success of any implicit sensor depends on two things: its ability to detect when it should be accepting data and how accurate that data actually is. Not surprisingly, as the implicit sensor becomes more strict with its acceptance criteria, the reported data becomes more accurate. The success of the Health Chair as an implicit sensor is therefore a combination of how often it accepts a rate and, of those accepted rates, how accurately it infers the rate compared to ground truth. Accuracy for a vital sign was calculated by comparing only those

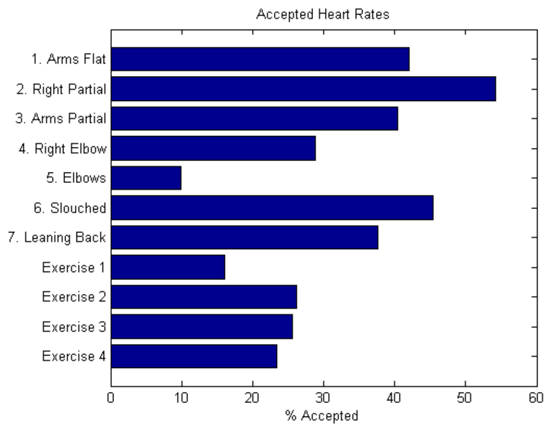

Figure 4. The percentage of accepted heart rates varies widely across the different positions occupants took. Positions with limited arm contact, such as elbows only in position 5, showed the least number of accepted heart rates.

rates found in a 10 second slide window that were accepted to the corresponding rate found in ground truth.

The laboratory study showed that the position an occupant is in has a large effect on the ability of the chair to be an implicit sensor. How often the Health Chair found valid rates differed greatly due to the position of the occupant. Overall, the Health Chair accepted heart rates 32% of the time and accepted respiratory rates 52% of the time.

Heart Rate

The detection of heart rate in EKG depends highly on the contact area of an individual's skin, the more contact the better. As shown in Figure 4 positions with less contact area between the participant's skin and the EKG electrodes accepted fewer heart rates. Most noticeably, the position that had the least amount of skin contact (elbows only) also had the lowest number of heart rates accepted.

Movement between the skin of the monitored individual and the electrodes also causes noise in the EKG signal. All the exercise experiments showed fewer detected heart rates, likely due to the movement of the participants after exertion. As a participant breaths heavily from the exercise, their body movements increase. This movement can cause motion artifacts in the EKG signal and, because the acceptance criteria for a rate are so strict, signals with these motion artifacts will often be rejected. The first exercise experiment's results show this to a greater degree than the others because many participants paced themselves after the initial exercise experiment.

Of the accepted heart rates, 83% were within 5 beats per minute of the ground truth heart rate. Figure 5 shows the accuracy of the Health Chair when compared to ground truth within 1 to 10 beats per minute. Unsurprisingly, with the low number of accepted rates detected for position 5, the accuracy of the heart rates in this position is high. Other positions, such as 3 and 4, show lower accuracy and a higher accepted percentage.

In the exercise experiments, heart rate was detected with 5 beats per minute of ground truth 75% of the time and within 10 beat per minute 85% of the time.

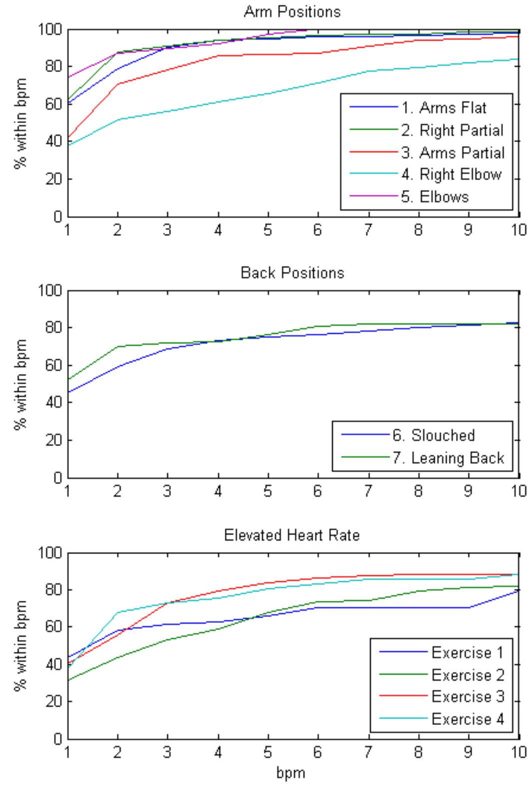

Figure 5. The accuracy of the accepted heart rates when compared to ground truth vary only moderately with position.

Respiratory Rate

Unsurprisingly, changes to the back positions had a large effect on the number of accepted respiratory rates as Figure 6 illustrates. For both positions, leaning forward (position 6) and leaning back (position 7), the percentage of accepted respiratory rates was less than half of the straight back position (position 1). In both cases this was due to the fact that fewer sensors were able to detect the respiratory motion of the upper body. When the participants leaned forward they came off the upper pressures sensors, effectively reducing the num-

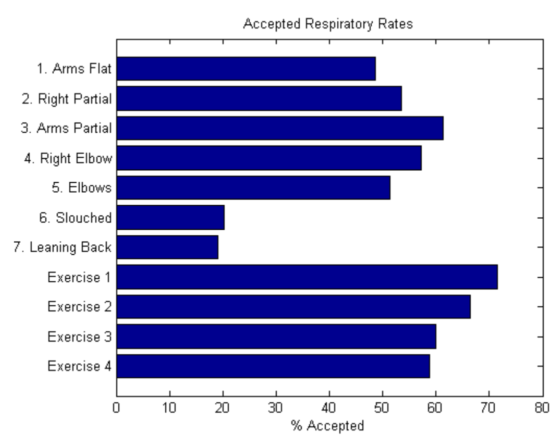

Figure 6. The percentage of accepted respiratory rates is fairly consistent across most positions, but drops sharply when the back position is changed in positions 6 and 7.

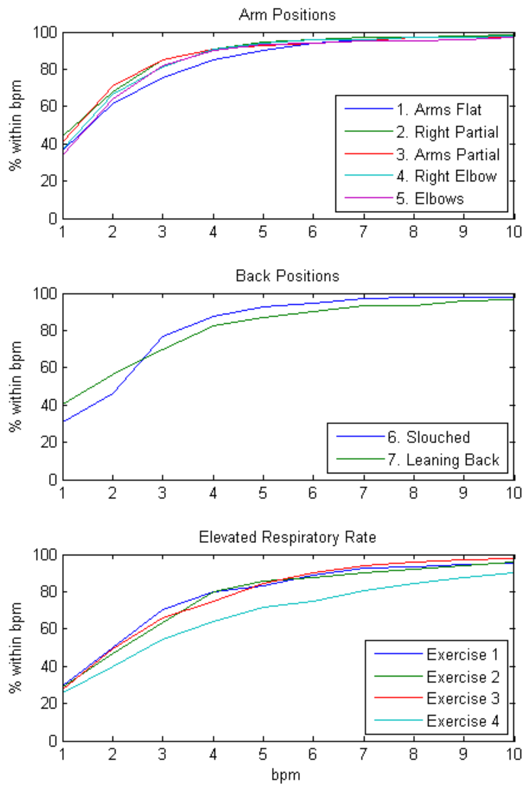

Figure 7. The accuracy of the accepted respiratory rates when compared to ground truth vary only slightly with position.

ber of sensors that could detected their respiratory rate. When participants leaned back they pressed themselves into the upper pressure sensors in such a way that the sensors reached their maximum pressure limit and therefore could no longer detect changes in respiratory motion.

Changing arm positions also had an effect on the respiratory rate detection. Position 3, where the participants had their upper forearms on the armrests and their hands in their laps, showed the highest percent of valid respiratory rates. This may be due to the fact that this position forces a slight slump to the upper body, pulling it slightly away from the sensors to prevent hitting the pressure limit, but not far enough that we see the drop in accuracy of position 6. As the participants move to having their elbows on the armrests, as in position 5, they are able to correct this slight slump and the percentage of accepted respiratory rates falls again.

As people exert themselves, they take deeper breaths, causing more expansion to the upper body. This can be seen in the accepted respiratory rates of the exercise experiments. With these deeper breaths the values from the pressure sensors change more drastically from an inhale to an exhale than seen from restful breathing. With this higher peak-to-peak amplitude, more rates are accepted by the respiratory rate signal processing. Opposite of the heart rate analysis, as the participants paced themselves through later exercise experiments this percentage of accepted respiratory rates fell.

Figure 7 shows the accuracy of the respiratory rate detection when compared to ground truth. Of the accepted respiratory rates, 73% were within 3 breaths per minute of ground truth.

REAL WORLD USAGE STUDY

To understand if the Health Chair would be effective as an implicit sensor in a real world environment we performed a series of in-situ studies. The goal of these studies was to provide an initial look into how an implicit sensing chair could be used in a real environment, and the frequency with which vital signs could be sensed. Because the intent of this real world study was to examine the chair in as natural an environment as possible we chose not to collect ground truth information, believing that wearing either of our ground truth devices might severely change the habits of our participants. For the in-situ study we focus not on the accuracy of the data collected, but on the participant's interaction with the chair.

Two Hour Study

Eight office workers at our company (M:4 F:4) participated in these studies. Each participant chosen was already using an identical, un-instrumented copy of the Health Chair as their normal office chair. Participants were encouraged to adjust the chair's arms, back, and height to match their normal usage. The health chair was then given to the participants to use in their offices for a period of two hours.

During the two hour in-situ period participants were asked to go about their normal work and act as if the Health Chair was their regular chair. We requested that they talk, type, read, write, and act as they normally do. Participants were asked to select a two hour period where they intended to sit for the full two hours, to ensure that we collected data that corresponded to an occupied chair. We did, however, provide them with a method to mark any times in that period where they stood up and left the chair if required. We also asked that participants work without sweatshirts or sweaters to make sure that their forearms were bare.

Given how often calculated rates are rejected as invalid even when the person is in full contact with the chair sensors, we expect an even smaller number of valid detected heart rates in real world usage, where an occupant is not always in contact with the backrest or armrests. Extrapolating from the lab study, where 32% of the time the chair accepts heart rate, we can assume that for a period of two hours if an occupant only uses both armrests for 10 minutes then we will only accept heart rates for 2.7% of that period (about 3 minutes).

Results from the two hour in-situ experiments, Table 2, show that while the chair does not pick up health signs for a large portion of natural use, it detects some resting heart rates and respiratory rates. Each detected rate in Table 2 corresponds to 10 seconds of vital sign signal from the occupant. As indicated by our usage study the use of the backrests, and therefore the detection of respiratory rate from the backrest, is more common than the use of the armrests. Additionally, most of the in-situ participants reported using the armrests very infrequently during the two hour period.

Participant	Resting Respiratory Rate Detected	Resting Heart Rate Detected
1	8.2% (9.7 minutes)	0.30% (21 seconds)
2	7.2% (8.7 minutes)	0.18% (13 seconds)
3	6.2% (7.2 minutes)	0.45% (33 seconds)
4	6.7% (8.2 minutes)	0.78% (57 seconds)
5	4.6% (5.5 minutes)	0.42% (30 seconds)
6	5.9% (7.1 minutes)	0.33% (24 seconds)
7	5.8% (7.1 minutes)	0.21% (16 seconds)
8	5.5% (6.7 minutes)	0.18% (13 seconds)

Table 2. Amount of time during two hour in-situ study that resting rates (40-100 beats per minute for heart rate, 12-20 breaths per minute for respiratory rate) are detected during normal chair usage.

Eight Hour Study

Due to the limited time heart rate was detected in the two hour in-situ studies we decided to conduct three longer in-situ studies. In this experiment we looked to evaluate the upper bound on sensing heart rate from the Health Chair. We recruited three additional participants who responded to the chair usage survey with ”often uses armrests”. These three office workers (M:1 F:2) were given the Health Chair to use as their normal chair for one workday, a period of about 8 hours. Similar to the two hour study, participants were asked to go about their normal work. For this study we put a pressure sensor in the seat of the chair to detect when the occupant was present so participants did not need to record their absences. As this study focused on heart rate, respiratory data was not collected.

To determine what role electrical or mains noise played in the lower numbers of detected heart rates of the in-situ studies, participants were first asked to sit in the chair with their arms on the armrests for 5 minutes in an area with minimal amounts of this noise. We then took the chair and participant to their office and repeated the five minute sitting period in the new environment where the 8 hour study would take place. Then the Health Chair was left in the office of the participant to use throughout the rest of the day.

Results from the eight hour in-situ experiments show that the different environments played a significant role in reducing the number of accepted heart rates (see Table 3). Heart rates could be detected in the minimal noise environment to a much greater degree for each participant than in their respective offices. Table 3 shows the percentage of time the participant was detected sitting during that 8 hour period and how often a resting heart rate was detected when they were sitting. In this longer study, heart rate was detected very infrequently; although the absolute amount of time with a valid heart rate was longer, the relative percentage of time heart rate was detected was similar to the shorter in-situ study.

BARRIERS FOR IMPLICIT SENSING IN THE WILD

Moving the Health Chair from the laboratory to the real world gave us insight into the pragmatic challenges of using implicit sensors in the wild. Though the primary goal of an implicit sensor is to be an opportunistic sensor, sensing only when the opportunity to collect data presents itself, we would still like to collect as much data as possible. When moving the Health Chair into a real world environment we saw a significant drop in the accepted respiratory and heart rates. From this experience, we describe a few of the barriers that we discovered to using the Health Chair as an implicit sensor in the wild.

	Low Noise	Office Noise	Sitting Time	Heart Rate
Length	5 min	5 min	8 hrs	8 hrs
1	55%	0.39%	63.89%	0.53% (1.8 min)
2	31%	4.5%	50.42%	0.37% (.88 min)
3	31%	16%	76.29%	0.68% (2.1 min)

Table 3. The first two columns show percentage of time heart rate was detect in a low noise environment and in the participant's office under ideal conditions. Heart rate was much harder to detect in the in-situ office environment likely due to electrical noise. The last two columns show the amount of time participants were sitting in the Health Chair over an 8 hour period and the percentage of time heart rate was detected while sitting.

Challenges

Usage Study ≠ Natural Usage

In the wild, people do not always use objects in the way one would expect. For the Health Chair we concluded originally from our usage study that the armrests were the ideal locations from which to sense heart rate. This location provided us with skin contact, sensor locations on either side of the heart, and was reported to be a location on the chair used by 80% of our respondents to some degree. However, heart rate was the least detected vital sign in our in-situ study - to a much lower degree than our laboratory study results would lead us to believe.

This lack of detected heart rates may be greatly attributed to participants in the in-situ study simply not using the armrests for the study period. Additionally, an occupant's use of the armrests may not be suited to our sensors. Using only one of the armrests of the chair is still 'use', and would have been a correct response to our usage study, but would not provide the Health Chair with the conditions necessary to sense heart rate. We suspect that many of the participants of the in-situ study exhibited some of these behaviors naturally and therefore contributed to the loss of sensing opportunities.

Work Stops For No Implicit Sensor

People don't just sit in chairs, they use them as a physical platform to work from. How people use a chair's armrests depends significantly on their current task. For example, when typing an occupant is much more likely to have their arms resting on the desk in front of them than on the armrests. When watching a video or reading something on a screen they may also lean forward onto the desk, or lean back and rest their arms on the armrests or in their lap. These different positions can greatly change what the Health Chair is able to sense. These tasks may be short or long and vary significantly across an occupant's day or week. Additionally, people with different occupations may perform tasks unsuited to the sensing capabilities of the Health Chair to a greater degree than others, such as programmers leaning on their desk or keyboard to type.

Even if the occupant is in contact with the implicit sensors the task they are performing can interfere with rate detection. If the occupant is talking to a colleague, as likely happened in our study due to a number of shared offices, then the pressure sensors may not be able to detect a sensible respiratory rate despite full contact with the occupant. If the occupant is reading a paper with both arms on the armrests, but marking up the paper with pen and flipping pages, the constant or near constant motion of their arms can make heart rate detection impossible. In some cases, the natural use of the chair may even trick the Health Chair into detecting a false rate in ways that we did not observe in the laboratory study. If a person is fidgeting at a regular rate or bobbing their head and body in time to some music, the Health Chair pressure sensors could pick up that rate instead of their respiratory rate.

Electrical Noise and Physical Environments Vary Widely

The physical environment of a lab is almost never the same as the real environment in which a system will be deployed. For EKG the real world has the added complications of being a very noisy environment for electrical signals. Due to the nature of EKG, if a individual puts their foot over a power supply the 60Hz noise can cause interference with the EKG signal. In the controlled environment of the lab, this noise can be missed because that particular environment may not have a problematic power source. In the wild, you cannot know what will be around a chair, especially if no action can be asked for and taken by the user to mitigate such noise. We observed this effect in the 8 hour in-situ study.

Indirect interference from the environment can also be a large problem, and much more difficult to identify. For example, the EKG electrodes require skin contact to pick up the electrical signals stimulating the heart. The more conductive an occupant's skin is, the easier it is to detect a clear signal. Since the skin's conductance is governed by the amount of sweat on a person's skin, in an environment that induces sweating, such as an office room that is warm, EKG detection will be easier. In this way the real environment affects the occupant, which affects the Health Chair's ability to detect vital signs. In a similar manner, a cold environment may induce the occupant to wear thicker, heavier clothes, thus making their respiratory movements less pronounced and harder to sense.

Future Work

The challenges found in our move from that lab to the wild informs the direction of future work, for both our Health Chair and other implicit sensors. It is important to remember that the goal of an implicit sensor is to be an opportunistic sensor. It is not designed to constantly detect someone, but it should take advantage of every opportunity to perform sensing. Future work in this area would both address the limitations of the sensing design of the Health Chair and examine other opportunities for implicit sensing available in the Health Chair's primary environment.

Future work on the sensing systems of our Health Chair, including both the sensor design and signal processing, would likely improve the percentage of accepted rates that can be detected and the accuracy of the accepted rates. Improvements such as detecting and removing motion artifacts in the respiratory signal processing, rather than rejecting any window showing motion, is a direction of future work. Since users make frequent movements in a chair, work on more advanced signal processing to discriminate between signal and noise may also help ease the strict acceptance criteria for EKG signals, allowing for heart rate to be measured more frequently. Additionally, since the current rejection algorithm for heart rate uses the standard deviation of the RR intervals measurements of some cardiac abnormalities, such as premature ventricular contraction and atrial fibrillation, may be rejected despite the signal displaying a true physical accurate heart beat. Future work on the heart rate detection algorithm would be needed to detect these signals from the Health Chair.

Additionally, alternate sensing systems, such as the through-clothes EKG being developed by Ford, may be able to further the sensing capabilities of the Health Chair by reducing some of the physical sensing limitations. Alternate sensing of respiratory rate could include performing impedance measurements across the chest or extracting respiratory sinus arrhythmia from EKG [5, 9]. Since the extraction of respiratory sinus arrhythmia can be done directly from the RR intervals of an EKG signal, we could condense the health sensing to only one sensor on the chair or compare the respiratory rate found in the EKG to the pressure sensors synchronously collecting respiratory rate to improve the accuracy of both. However, given our results using EKG in-situ we did not pursue this approach.

Our in-situ study showed us that the armrests of a chair is not always the best location to sense heart rate. The Health Chair's EKG sensors require skin contact from both arms on both armrests, behavior that our in-situ tests show may not be common. Additionally, since contact with skin is required, jackets, long sleeves and other arm coverings may interfere with the collection of EKG. These results indicate that the armrests may not be the optimal choice for EKG, often because the occupants are using the environment surrounding the chair to perform certain tasks. Therefore instrumenting the surrounding environment, such as the desk, keyboard, or mouse, may provide a better implicit sensor for sensing systems that require skin contact. Future work would examine these locations, both in concert with the Health Chair and on their own, as implicit sensors designed for use in the wild.

CONCLUSION

In this work we have shown that the Health Chair can accurately detect respiratory and heart rates in the laboratory though a variety of different positions an occupant can take. In the wild, we have shown that the Health Chair can detect respiratory and heart rates as the occupant uses the chair in a natural fashion, without active participation. Our chair usage study suggests that the Health Chair as an implicit sensor would be able to provide a large portion of the population with frequent measurements of their vital signs. The Health Chair is an implicit sensor that can opportunistically provide more frequent information about a person's vitals than they would be willing or able to obtain otherwise.

ACKNOWLEDGMENTS

We thank our participants, reviewers, Tom Blank, Dan Morris, and Asta Roseway for their assistance.

REFERENCES

1. Anttonen, J., and Surakka, V. Emotions and Heart Rate While Sitting on a Chair. *Proceeding of the SIGCHI Conference on Human Factors in Computing Systems* (2005), 491–499.
2. Arnrich, B., Setz, C., Marca, R. L., Troster, G., and Ehlert, U. What Does Your Chair Know About Your Stress Level? *IEEE Transactions on Information Technology in Biomedicine 14*, 2 (2010), 207–214.
3. Bu, N., Ueno, N., and Fukuda, O. Monitoring of Respiration and Heartbeat during Sleep using a Flexible Piezoelectric Film Sensor and Empirical Mode Decomposition. *Proceedings of the 29th Annual International Conference of IEEE EMBS* (2007).
4. Cheng, J., Zhou, B., Sundholm, M., and Lukowicz, P. Smart Chair: What Can Simple Pressure Sensors under the Chairs' Legs Tell Us about User Activity? *UBICOMM13: The Seventh International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies* (2013).
5. Ernst, J., Litvack, D., Lozano, D., Cacioppo, J., and Berntson, G. Impedance Pneumography: Noise as Signal in Impedance Cardiograhy. *Psychophysiology 36*, 3 (1999), 333–338.
6. Fletcher, R. R., Poh, M.-Z., and Eydgahi, H. Wearable Sensors: Opportunities and Challenges for Los-Cost Health Care. *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE* (2010), 1763–1766.
7. Ford Research Unveils Heart Rate Monitor Seat, 2013. `http://mobihealthnews.com/11014/ford-research-unveils-heart-rate-monitor-seats`.
8. Hamilton, M. T., Healy, G. N., Dunstan, D. W., Zderic, T. W., and Owen, N. Too Little Exercise and Too much Sitting: Inactivity Physiology and the Need for New Recommendations on Sedentary Behavior. *Current Cardiovascular Risk Reports 2*, 4 (2008), 292–298.
9. Hirsch, J. A., and Bishop, B. Respiratory Sinus Arrhythmia in Humans: How Breathing Pattern Modulates Heart Rate. *American Journal of Physiology - Heart and Circulatory Physiology 241*, H620-H629 (1981).
10. Hurst, A., Zimmerman, J., Atkeson, C., and Forlizzi, J. The Sense Lounger: Establishing a UbiComp Beachhead in Elder's Homes. *CHI'05 Extended Abstracts on Human Factors in Computing Systems* (2005), 1467–1470.
11. Kortelainen, J. M., Mendez, M. O., Bianchi, A. M., Matteucci, M., and Cerutti, S. Sleep Staging Based on Signals Acquired Through Bed Sensor. *IEEE Transactions on Information Tachnology in Biomedicine 14*, 3 (2010), 776–785.
12. Martins, L., Lucena, R., Belo, J., Santos, M., Quaresma, C., Jesus, A. P., and Vieira, P. inteligent Chair Sensor. *Communications in Computer and Information Science 383* (2013), 182–191.
13. Meyer, J., Arnrich, B., Schumm, J., and Tröster, G. Design and Modeling of a Textile Pressure Sensor for Sitting Posture Classification. *Sensors Journal, IEEE 10*, 8 (2010), 1391–1398.
14. Mota, S., and Picard, R. W. Automated Posture Analysis for Detecting Learner's Interest Level. *2003 Conference on Computer Vision and Pattern recognition Workshop* (2003).
15. Mutlu, B., Krause, A., Forlizzi, J., Guestrin, C., and Hodgins, J. Robust, low-cost, non-intrusive sensing and recognition of seated postures. *20th annual ACM symposium on User interface software and technology* (2007).
16. Pantelopoulos, A., and Bourbakis, N. G. A Survey on Wearable Sensor-Based Systems for Health Moniotoring and Prognosis. *IEEE Transactions on Systems, Man, and Cybernetics* (2010), 1–12.
17. Poh, M.-Z., McDuff, D. J., and Picard, R. W. Non-contact, automated cardiac pulse measurents using video imaging and blind source separation. *Optics Express* (2010), 10762–74.
18. Postolache, O., Girao, P. S., Postolache, G., and Pereira, M. Vital Signs Monitoring System Based on EMFi Sensors and Wavelet Analysis. *Instrumentation and Measurement Technology Conference Proceedings* (2007), 1–4.
19. ECG processing - R-peaks Detection, 2011. `http://www.librow.com/cases/case-2`.
20. Rashvand, H., Salcedo, V. T., Sanchez, E. M., and Iliescu, D. Ubiquitous wireless telemedicine. *IET Communications 2*, 2 (2008), 237–254.
21. Stanford, V. Using Pervasive Computing to Deliver Elder Care. *Pervasive Computing 1*, 1 (2002), 10–13.
22. Tan, H. Z., Lu, I., and Pentland, A. The Chair as a Novel Haptic User Interface. *Proceedings on the Workshop on Perceptual User Interfaces* (1997), 56–57.
23. Tan, H. Z., Slivovsky, L. A., and Pentland, A. A Sensing Chair Using Pressure Distribution Sensors. *IEEE/ASME Transactions on Mechatronics 6*, 3 (2001), 261–268.
24. Tanaka, S., Nogawa, M., and Yamakoshi, K. Fully Automatic System for Monitoring Blood Pressure from a Toilet-Seat Using the Volume-Oscillometric Method. *Engineering in Medicine and Biology Society, 27th Annual International Conference* (2006), 3939–3941.
25. Toyota Demos ECG Sensing Steering Wheel, 2011. `http://mobihealthnews.com/12128/toyota-demos-ecg-sensing-steering-wheel/`.
26. Zheng, Y. J., and Morell, J. B. A Vibrotactile Feedback Approach to Posture Guidance. *Haptics Symposium, 2010 IEEE* (2010), 351–358.