PostureChair: A Real-Time, As-Needed Feedback System for Improving the Sitting Posture of Office Workers

by

Jessica Speir

A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs in partial fulfillment of the requirements for the degree of

Master of Computer Science

in

Human-Computer Interaction

Carleton University
Ottawa, Ontario

Abstract

Computer users spend the vast majority of their time sitting, and poor posture in the workplace is an endemic issue. This work presents PostureChair, a posture detection system that uses contextual digital feedback to persuade users to improve their sitting posture. Two types of digital feedback, with varying amounts of information, were compared through a repeated measures study to determine how much information is necessary to improve posture and to appeal to the user. The results of the study showed participants' sitting posture improved significantly with both feedback types when compared to their posture with feedback disabled. Participants overwhelmingly preferred the more detailed feedback even though it did not clearly improve users' sitting posture beyond the simpler feedback. The PostureChair system was well received and demonstrates that contextual posture improvement is an effective and much-needed addition to the workplace.

Acknowledgements

I would like to thank my research supervisor Anthony Whitehead for providing me with this opportunity, and for keeping me on track and on schedule.

To my comrades-in-arms Rufino Ansara, Shannon Fenton and Colin Killby, you kept me sane when things were tough, gave me a hand when I needed it, and provided distractions at just the right time. Thank you for your friendship.

My heartfelt gratitude goes to my editors, Heather Stanger, Dennis Arsenault and Marjorie Speir. Any mistakes that remain in this work are mine alone. Dennis, you have been my rock throughout this project. You led the way, and I absolutely could not have done this without you.

Thanks also to my parents Richard and Marjorie Speir; you put me on this path and have been unwavering in your support my entire life. You gave me the confidence to push myself and never stop trying.

Finally, I want to dedicate this work to my grandfather, Don Speir, who passed away before it was completed. He believed strongly in the value of education, and he was so proud of what his children and grandchildren accomplished. This is for him.

Table of Contents

List of Tables

List of Figures

Chapter 1 Introduction

Computer use has increased dramatically over the past few decades. In a 2010 time-use survey, the most recent data available, nearly five times more Canadians reported using computers during their leisure time than in 1998 [54]. On average they spent 83 minutes per day using a computer for leisure. A majority of Canadian workers also use a computer at work [34]. Nearly all of this computer use is done sitting down, despite the recent development of smartphones and tablets that conceivably allow for computer use in almost any position. A person in the developed world who works primarily at a computer could reasonably sit for up to 15 hours a day [25]. It is easy to imagine a typical office worker's day, commuting to work by public transit or driving, working at their computer for 8 hours, with an hour for lunch (seated in the lunch room, of course), the same commute home, and an evening on the couch in front of a TV. This hypothetical worker might be able to fit in 30 minutes of vigorous exercise before or after work, but that does little to change the overall sedentary pattern of their day. High rates of computer use, especially at work, are an important factor in the sedentary lifestyle that most Canadians experience today.

All of this sitting comes with significant health costs, both from inactivity and from poor posture. Long periods of sitting have been linked to obesity, cardiovascular disease and premature mortality, although there is also evidence that these adverse effects can be mitigated by short standing breaks [46]. Poor sitting posture has been identified as a risk factor for musculoskeletal disorders [58], and particularly for lower back pain [23]. Musculoskeletal disorders can cause chronic pain in the limbs, neck and back. Preventing pain and improving the quality of life of computer users is a valuable goal on its own, but chronic pain is also an expensive problem for employers and the health care system. For instance, back pain alone made up 18% of lost time claims in Canada in 2013 [60]. In recent years various workplace innovations have been introduced to reduce sitting, such as standing or treadmill desks, and standing meetings. However it is likely that prolonged periods of sitting will continue to be common in computer-based workplaces for the foreseeable future. Instead, improving computer users' sitting posture is crucial to reducing the costs on individuals, employers and the health care system.

Many attempts have been made by governments and employers to educate computer users on the risks of poor sitting posture and the ideal, or ergonomic, postures they should adopt. These ergonomic interventions have had some success [2, 11] however poor posture is still an endemic issue in the workplace. Sitting posture is a continuous physical state that lasts over many hours and normally requires little to no attention, unless it causes immediate pain or discomfort. Computer work, on the other hand, requires most or all of an individual's attention. For a normal computer user, continuously adjusting their sitting posture while maintaining focus on their work is nearly impossible without the immediate feedback provided by pain, at which point the damage has already occurred. Changing habitual behaviours like sitting posture requires feedback, and since the human body does not supply sufficient feedback until it is too late we must turn to technology to provide the solution.

1.1 Posture Detection

There are many examples of persuasive technologies designed to detect and modify computer users' sitting posture in academic literature and, in very recent years, commercially (e.g. [6, 33, 47]). Posture detection can now be done fairly reliably, although because of the variability in the human body most works rely on calibration to known users. Detection methods include video-based systems (e.g. [27, 47, 56]), wearable sensors (e.g. [14, 51]), and pressure or force sensors incorporated into office chairs (e.g. [10, 24, 37, 52, 55, 62]). Posture detection using an office chair augmented with sensors is the most prevalent method in the literature.

1.2 Posture Modification

There also has been considerable academic work done on the methods of feedback that can be used to modify computer users' posture. These feedback methods are usually divided into three categories: physical items, digital feedback and vibrotactile (i.e. vibration) feedback. In this context the term "digital" feedback is used to describe any feedback that appears on a screen, as opposed to feedback that is simply accomplished with computers. Feedback that is described as "physical" or "vibrotactile" in the literature can be accomplished digitally, but only the on-screen feedback is categorized as "digital". Several of these works use the posture detected through an augmented office chair to trigger the appearance of the feedback and then examine the effect on the users' posture (e.g. [10, 13, 21, 62]). The advantage of combining posture detection and feedback is that the feedback can be provided on an as-needed basis, so that it is only given when the computer user has slipped into bad posture. This contextual aspect reduces the possibility that the feedback will be ignored as background "noise" since it will bring the user's attention to their posture only when necessary. While there are works that investigate the effectiveness of vibrotactile [13, 21, 62] and physical [10] feedback, the effectiveness of digital feedback in this type of system has been neglected. Thus there is still work to be done investigating whether digital feedback is actually effective at modifying computer users' sitting posture, as part of a posture detection system.

1.3 Research Questions

Our first research question attempts to address the gap in the literature regarding how effective digital feedback is at modifying computer users' posture in a context-aware (i.e. posture-detecting) system:

Research Question #1: *Can contextual, digital feedback improve the sitting posture of computer users?*

To investigate this question we developed a posture detection and modification system called PostureChair. The primary purpose of this system is to examine the effectiveness of the feedback, so the method of posture detection was based on the most commonly accepted model found in the literature (an office chair augmented with force sensors), with few changes. As specified in the research question, the PostureChair system attempts to improve the user's posture through persuasive, digital feedback. A within-subject user study was conducted to address the research question. Participants in this study used the PostureChair system for a single period of 1.5-2 hours, divided into three sessions, one session without feedback and two sessions with feedback. Participants' detected posture was then compared between the feedback-enabled and feedback-disabled sessions in each study.

Unlike physical or vibrotactile feedback, digital feedback is very flexible in the amount of information that can be presented to the user. This information can range from simple binary feedback (e.g. "You have good/bad posture") to more complex instructions on how to improve posture and information on the risks of poor posture. Our second research question addresses the amount of information required to make the digital feedback effective:

Research Question #2: *Does digital feedback need to include specific instructions on how to improve posture in order to be effective?*

The PostureChair system includes two different types of digital feedback: Generic and Specific. Both types of feedback contain a simple text notification of poor posture, but the Specific feedback also includes information on how the user can improve their posture through a real-time, abstract representation of their current posture. The user study included two feedback-enabled sessions, one for each of the feedback types. Our second research question was addressed by comparing the detected posture in the two feedback-enabled sessions with each other, as well as against the feedback-disabled session.

1.4 Contribution

Our work on the PostureChair system joins a growing body of literature in posture detection and modification. It uses a well-established posture detection method to provide context for digital feedback. This combination has been used in prior works however this is among the first to evaluate how effective it is at actually improving computer users' posture.

The participants' quality of experience with digital feedback is also examined in this work. Any commercial persuasive system must take into account how well users will accept the feedback. If users find the feedback irritating they will soon abandon the system, regardless of how effective it is. These aspects of the digital feedback are examined through a questionnaire administered to the participants at the end of the user study and through informal observations of and conversations with the participants.

Finally, this work provides additional contributions through incidental results related to sitting break times, participants' general knowledge of ergonomics, and their overall perceptions of the PostureChair system.

1.5 Thesis Overview

We begin the body of this thesis by discussing background topics and notable related works in Chapter 2. This includes defining and describing persuasive technology to give context to this work, identifying the opportunities for behaviour detection and modification in the typical computer workstation, and describing the currently accepted recommendations for sitting times and good sitting posture. Related works are discussed according to their posture detection and feedback methods.

Chapter 3 examines the PostureChair system itself, starting with the design of the system in the context of the related works and other design factors. This Chapter also includes descriptions of the hardware and software components of the PostureChair system. We then proceed to outline the user study that examined the effectiveness of the system in Chapter 4, including details on recruitment and the data collected.

The results of this study are presented in Chapter 5, beginning with a description of a code error that was discovered during the user study, how it affected the data and some unexpected results that were found as a result. The data related to whether or not this system is effective at modifying a user's sitting posture is then described and analyzed. This is followed by data and analysis related to the participants' experience of the digital feedback, and finally a few incidental results on posture knowledge and breaks in prolonged sitting. The implications of the data and any limitations are discussed throughout this chapter. Finally, the findings of the thesis, and possible directions for future work, are summarized in Chapter 6.

Chapter 2 Background

2.1 Persuasive Technology

Persuasive technology includes “any interactive computer system designed to change people’s attitudes or behaviours” [19]. The best way to understand what falls under the umbrella of persuasive technology is to consider what is excluded from this category. Importantly, *persuasion* in this context excludes any attempts to create change through deception or coercion [41]. The latter involves the use of force or threats. To be considered persuasive the user’s change in attitude or behaviour (or both) must be entirely voluntary. Also excluded from persuasive technology is any technology that only mediates communication between humans; the interaction that leads to the change in behaviour must be between the user and the computer, not between two users. For instance, if one person attempts to persuade another to quit smoking through instant-messaging, the instant messaging software is simply the medium of communication and would not be considered persuasive technology. Finally, persuasive technology includes only technology that is designed and planned to be persuasive. Any technology that has an unintended side-effect of changing users’ behaviour or attitudes is excluded [19].

2.1.1 Persuasive Technology by Method of Persuasion

When examining persuasive technologies it is important to consider how the persuasion can be best accomplished. According to BJ Fogg’s seminal book “Persuasive Technology”, these technologies can be categorized according to seven different methods of persuasion: Reduction, Tunneling, Tailoring, Suggestion, Self-Monitoring, Surveillance and Conditioning [19]. Often a particular technology will use methods from multiple categories.

Reduction technologies reduce a complex task into a few simple steps to persuade users to take on the complex task. A common example is the "one-click" purchase feature on Amazon.com, which remembers shipping and billing details so customers do not have to input them repeatedly.

Tunneling technologies reduce the choices available into one, easy-to-follow path, as in the case of software installation wizards that walk you through each step of that process.

Tailoring technologies customize the information presented to the context, the user, or both. Users are more likely to respond to information that they feel is tailored to their situation. Google Ads that reflect a user's Internet browsing history fall into this category.

Suggestion technologies suggest a new behaviour at the most opportune moment. For instance, the small digital sign at the side of the road that displays your current driving speed can be classified as a suggestion technology because of the location of the display and the accompanying sign showing the actual speed limit. It can also be classified as a tailoring technology because it displays your current driving speed.

Self-monitoring technologies present information like heart rate or steps taken in a day in real time so that users can see how close they are to a particular goal. These technologies are often used for healthcare-related goals.

Surveillance technologies involve overt, outside monitoring of the user's behaviour. For instance, a computer worker who is aware that their employer monitors their Internet browsing may change which sites they visit. This can be considered persuasive technology only if the employer intends to use the surveillance to persuade their employees to stay away from time-wasting websites.

Finally, *Conditioning* technologies involve offering rewards in exchange for a change in behaviour or attitude. This technique is common in computer games, which offer in-game rewards for continuing to play the game.

2.1.2 Persuasive Technology by Type of Change and Outcome

Another way to categorize persuasive technology is to look at what the technology is attempting to change in the user, and what has to happen to achieve that change. Oinas-Kukkonen's model categorizes persuasive technology along these lines [41]. In this model, persuasive technology attempts one of three changes: ensuring that the user complies with the desired behaviour (C-Change), causing a more enduring behavioural change (B-Change), or changing the user's attitude (A-Change). C-Changes are the easiest to achieve, while A-Changes are the most difficult and much less studied. In order to achieve any of these changes one of three outcomes must occur, according to Oinas-Kukkonen's model [41]: a new behaviour or attitude must be formed that did not exist before (F-Outcome), the user's response to an issue or trigger must be altered (A-Outcome), or the user's current attitude or behaviours must be reinforced and made more resistant to change (R-Outcome).

2.2 Changing Computer Habits

A majority of Canadian workers now use a computer at work [34], and even more use a computer in their leisure hours [54]. This amount of time spent on a single activity makes the habits that workers develop while using a computer extremely influential on their overall health and well-being. Persuasive technologies that aim to improve people's health and well-being have an advantage in that users of these systems are often intrinsically inclined to be persuaded towards these goals [40], which can make behaviour or attitude change much easier to achieve. Bringing persuasive technology into the computer workstation is a promising avenue for influencing the habits crucial to computer users' health. This section includes descriptions of the methods by which feedback can be communicated to the user, as well as the avenues of behaviour detection available in the typical computer workstation. This is followed by descriptions of persuasive commercial and academic works that involve changing computer habits not related to posture. Relevant posture-related works are described in Sections 2.4 and 2.5.

2.2.1 Feedback

Many of the categories of persuasive technologies rely on the detection of certain behaviours followed by the communication of that information back to the user. This communication is commonly referred to as feedback. Several of the categories in BJ Fogg's model described above rely on feedback [19]. Tailoring systems use the detected behaviour, context, or user to determine what information is presented in the feedback. Suggestion systems attempt to give feedback at the most opportune moment. Self-monitoring systems communicate the detected information immediately back to the user, and Surveillance systems use feedback to ensure the user is aware that their behaviour is being detected. The feedback itself is very important to the success or failure of these types of persuasive systems.

Feedback is usually described by the modality, or "channel", by which the information is communicated to the user. Common categories of feedback related to computer behaviour and posture include:

- *Digital* feedback, which is presented on a screen of some kind, usually a computer monitor
- *Physical* feedback, in which information is communicated by a physical agent or item, usually on the desk surface
- *Vibration/Vibrotactile* feedback, usually through the office chair as it is the only surface we can be sure the user is in constant contact with
- *Audible* feedback, which is usually paired with another modality as an extra measure to bring attention to the feedback

Although all of these feedback modalities can be accomplished digitally, through the use of computing technology, only feedback that occurs on a screen is labelled as Digital. Digital feedback could also be called visual or graphical feedback, however digital is the term commonly used in the literature and so it is the term used throughout this work.

Each of these modalities comes with advantages and disadvantages depending on the workplace environment and the amount of disruption to their workflow the user finds acceptable. For instance, audible feedback can be very effective at commanding the user's attention, but may not be appropriate in a shared workspace. Physical feedback can be novel and often incorporates affective elements, but it might be more easily overlooked than the other modalities. Digital feedback is the most familiar to users and is often the easiest to implement, but it can be very disruptive to workflow. Vibrotactile feedback can be much less disruptive to workflow since it uses the sense of touch, which is not usually involved in computer work; however vibrations can be irritating and may alienate some users. These advantages and disadvantages must be considered and balanced as much as possible when designing a persuasive system for the workplace.

2.2.2 The Typical Workstation

Although the computer is an important component of any office-related persuasive technology, we can look beyond simple software to other items in the typical workstation for opportunities for behaviour detection and feedback. First, we must consider what makes up a typical computer workstation. There is likely a desk, perhaps with personal items on display such as framed photos and small potted plants, and an office chair with some degree of adjustability. The computer may be on the desk or underneath it, with the monitor on the desk. The monitor may or may not include a built-in camera and is hopefully adjustable in height. There is a keyboard and a mouse, ideally on a tray at a different level than the desk for maximum adjustability. Finally there are the peripherals, perhaps a printer and most likely a smartphone or tablet. Each of these items is an opportunity to detect a particular behaviour or to give feedback that may lead to a modification of that behaviour. Many of these opportunities have been explored in academic literature and in commercial products.

2.2.2.1 Software in the Workstation

The earliest persuasive technologies in the typical computer workstation were software-only products, measuring mouse and keyboard actions and providing only digital feedback. One such commercial product is RSIGuard [48], which models user behaviour based on their mouse and keyboard use and has a wide variety of features aimed at reducing muscle strain and other repeated-stress injuries. One feature included in RSIGuard suggests long breaks for the user to rest away from the computer and "microbreaks" to promote self-awareness, based on context and user preferences. There is also a tool that gives tailored information on applying ergonomics to specific workstation setups, and another that implements hot keys and auto clicking to replace high strain activities (e.g. mouse clicking) with lower strain activities. RSIGuard can even limit computer use entirely by forcing the user to log off or shut down, although this last feature involves some level of coercion and cannot be considered entirely persuasive.

An entirely software-based product in academic literature is the work of Berque et al. [7], which is aimed at improving typing behaviour. Berque et al. divide typing behaviour into three different components, which they address individually: typing speed, use of shortcuts, and taking breaks. Each of these behaviours is addressed with a different method of persuasion through digital feedback. Several of these methods incorporate affective or emotional attributes: a small smiley face in the system tray changes from a smile to a frown as the user's typing speed increases, pop-ups indicate when it is time to take a break and when an opportunity to take a shortcut has been missed, and users are persuaded to take breaks by providing fun activities during the break. This last method is an example of a conditioning persuasive technology, since the fun activity is a reward for taking the break. The system also keeps track of the user's word usage in order to suggest suitable keyboard shortcuts, a tailoring feature. Although only the typing speed and shortcut creation were tested in Berque et al.'s work, they did find that the digital feedback was effective at persuading users to choose healthier behaviours.

2.2.2.2 Sensors in the Workstation

To move beyond simple keyboard and mouse usage requires incorporating sensors into other components of the workstation. The resulting persuasive technologies combine sensor-based behaviour detection with one of the feedback modalities described in Section 2.2.1. The most straightforward route for behaviour detection is to augment physical components that already exist in the workstation with proximity or pressure sensors. The Habit-Aware Mouse of Sonne and Grønbæk [53] uses conductive ink to form ultra thin capacitive sensors on the surface of the mouse buttons that can detect how close the user's fingers are to the surface of the mouse. Their aim was to prevent users from hovering their fingers just above the mouse for long periods of time, which can cause forearm pain. Sonne and Grønbæk did not implement any feedback or persuasive components to their system but explicitly stated those as intended future work. The most commonly augmented surface in the computer workstation is the office chair; the works in this area are explored in more detail in Section 2.4.3.

There are other sensors that can be leveraged for behaviour detection in the typical workstation, although they are used less often in the literature related to posture detection. Many monitors now have built-in webcams, so video-based detection of computer behaviours is possible. Cameras can also be added to the workstation, although this is more invasive than using built-in cameras. Added cameras can introduce more sophisticated detection capabilities such as eye-tracking or skeleton tracking. Video-based posture detection is addressed in Section 2.4.1. Smartphones are a new avenue for behaviour detection and feedback in the workplace that have the advantage of multiple built-in sensors and portability between workstations. These advantages have been leveraged by a variety of systems designed for general health interventions not limited to the workplace, as described in Klasnja and Pratt's review [30]. However, there is no guarantee that every office worker will have a smartphone, and the differences in capabilities between smartphone models can be a significant downside. Finally, wearable technology can also be introduced into the workplace, and can provide extremely customized detection for a wide variety of behaviours. Recent popular adoption of commercial products like activity monitors and smart watches are increasing the number of wearable sensors that users themselves are bringing into the workplace. However, wearables that can be used for posture detection often require skin-tight placement, which can be even more of an intrusion than cameras and can require a significant effort on the part of the user. Posture-related wearable systems are discussed more in Section 2.4.2.

2.3 Defining Good Posture

The first step to persuade users to adopt a better posture is to define good sitting posture. Only then can measurements of actual sitting posture be compared to this ideal standard. The following sections outline the recommendations for good sitting posture and for maximum sitting times that are currently generally accepted, in the context of the academic works that support them.

2.3.1 Current Recommendations for Posture

There are a wide variety of publications aimed at describing an “ergonomically correct” seated posture for office workers, complete with detailed instructions on how to achieve this posture [1, 9, 42, 44, 59, 61]. There are also online tools that allow users to customize these instructions to their height and desk setup, for example Ergotron’s Ergonomic Workspace Planner [16]. This ergonomic posture is widely accepted despite the paucity of academic studies to support its clinical advantage. There has been at least one attempt [8] to find quantitative evidence of the ergonomically ideal sitting posture, however it appears that there is too much variation across individuals to systematically define and demonstrate an ideal sitting posture for everyone. Sitting posture also involves complex full-body interactions, which makes this task even more difficult. The wide acceptance of the ergonomic posture, despite the lack of clinical evidence, is likely because ergonomic interventions in the workplace have been shown to have a significant positive effect on musculoskeletal symptoms [2] and worker productivity [11]. The generally recommended ergonomic sitting posture for office workers can be summarized with the following guidelines, taken from Canadian federal and provincial publications on office ergonomics [9, 61], and followed to some extent by other studies on posture [23].

1. Neck straight, head centred over the shoulders
2. Elbows at about 90°, arms at the sides
3. Hands in line with forearms, wrists straight
4. Back fully supported by the backrest of the chair
5. Thighs roughly parallel to the floor, with a 90°-120° angle at the hip
6. Knees at an angle of 90°-130°
7. Feet flat on the floor and fully supported

For this posture to be achieved by most computer users the chair seat should be roughly at knee height and the top of the computer screen should be about level with the user’s eyes when seated. The keyboard and mouse should be at the same height as the armrests of the chair, to allow the armrests to take the weight of the arms at the elbows. The backrest of the chair should have a curved lumbar support to follow the line of the back.

2.3.2 Current Recommendations for Sitting Time

As with the posture guidelines, maximum sitting time guidelines seem to be based on the professional judgment of clinicians rather than on quantifiable scientific data [50].

Popular recommendations include maximum sitting times of 20 minutes [50], 30 minutes [5] or a 5-minute break for every 40-50 minutes [43] or hour [45] of sitting. Sitting time guidelines are just now making their way into the workplace, and are much less established than ergonomic workstation recommendations. Current adherence to these guidelines is, as far as we know, abysmal. In a study of workplace adherence to sitting guidelines performed by Ryan et al., none of the participants met the 20 or 30 minutes maximum sitting time recommendations on the working days assessed [50]. However, the same study found that the 5-minute break per hour was far more achievable, even without educating the participants on sitting time recommendations. The same guidelines studied by Ryan et al. were also examined by Netten et al. [38]. The workers they studied spent 65-70% of their work day sitting, and on any given day only 5% were able to achieve the 20 minutes sitting guideline. By contrast 85% of the workers studied were able to achieve the 5 minute break per hour guideline. Taken together, these two studies show that persuasive technologies are likely necessary to encourage sitting times of less than 55 minutes. That many office workers spend the predominant part of their day sitting is indisputable; of the 550 people surveyed by Griffiths and Saponas, more than half reported sitting more than 9 hours a day [22]. The same study found that 91% of those studied had a “primary” chair in which they spent most of their time, suggesting that targeting a single chair for augmentation would be beneficial for a large number of office workers.

The evidence that sitting breaks are necessary is convincing, regardless of the confusion over what maximum sitting times should be recommended to the general population. Higher numbers of breaks in sitting time has been shown to improve measurements of metabolic health, independent of total sitting time or the amount or intensity of exercise performed [46]. The same technology that can be used to detect and influence posture can also be used to measure and influence sitting breaks. Taken together, posture and sitting breaks are two aspects of the same goal: to reduce the impact of sitting on the health and well-being of office workers.

2.4 Detecting Posture

As discussed in Section 2.2.2, the typical computer workstation contains several avenues for posture detection and for providing feedback to the user. A review of the literature showed three primary methods for detecting posture: cameras, wearable sensors, and instrumented chairs. This Section contains descriptions of related works in this field, discussed according to their method of posture-detection. Methods of feedback are discussed in Section 2.5.

2.4.1 Video-Based Systems

Human expert and non-expert judgement of a sitting posture is done visually, which makes a video-based posture detection system intuitive to many developers. The work of van Niekerk et al. [39] found that side-view photos of seated adolescents could produce valid and reliable posture indicators of the underlying spinal posture, using radiographs as the measurement of the ground truth. Their work mitigates any concerns that the external visual appearance might not reliably reflect the internal structure of the human body, and supports the use of video- and photo-based posture detection in other systems.

Many video-based posture detection systems take advantage of the Microsoft Kinect [29], a motion-sensing peripheral commonly used for video gaming that includes a skeletal estimation feature. The skeletal estimation was initially designed only for standing individuals, but was then extended for seated individuals whose lower body is not entirely visible to the sensor. This feature can be adapted for posture detection fairly easily however the Kinect still suffers from the disruption of adding a new peripheral to an existing computer workstation. In addition, the Kinect skeletal detection requires at least 80cm between the user and the camera [36], which may not be achievable in the limited space available in most workstations.

The Kinect has been used successfully by Taylor et al. [56] to detect standing posture as part of their Augmented Reality Posture Mirror (ARM). That system uses the Kinect to project the image of the user onto a screen in front of them, mimicking the effect of a mirror. Another system that used the Kinect successfully was BITAIKA of Ishimatsu and Ueoka [27]. Due to the limitations of the front-view seated skeletal detection Ishimatsu and Ueoka chose to place the Kinect side-on to the user. They combined this method of detection with an instrumented office chair to take advantage of the strengths of both methods. However, adding a side-view camera to a computer workstation is even more disruptive than adding a front-view camera. Both these systems show that video-based posture detection using the Kinect has potential, but also that there are serious limitations that must be overcome before they become practical in a real-world office setting.

We were only able to discover a single commercial solution that used a camera to detect sitting posture: the Philips ErgoSense Monitor [47]. Very little information is available about this product however it appears that the monitor has a built-in camera that it uses to detect the user’s head angle and distance from the screen. Although these are important aspects of sitting posture and are often cited in the overall ergonomic setup of a computer workstation, they do not describe the user's entire sitting posture.

2.4.2 Wearables

Wearable posture-detection systems are less common than video or augmented chair systems in the literature, possibly because they are more inconvenient for the user, often requiring sensors to be applied directly to the skin or special clothing. However, by measuring directly on the body there is often a higher confidence in the reliability of the results, as they are less affected by interference from the user's clothing and other visual barriers that often affect video-based system. One example of posture-detection with wearables is the work of Dunne et al. [14], which uses fibre-optic bend sensors worn along the spine to measure seated spinal posture. In this work the results of the bend sensors were found to be as accurate and reliable as expert visual analysis. The bend sensors were applied to the participants' skin or to skin-tight clothing, and the participants were seated in a backless chair to achieve the most favourable result. This requirement for the sensors to be as close to the skin as possible is a serious disadvantage of this method. It is also likely that in an office setting the back of the chair would interfere with the sensors or create discomfort for the user.

Wearables might be a more preferable posture-detection method when the target user base is more mobile than the typical office worker. Children are far more mobile than office workers, and would likely be more attracted to the novelty of a wearable robot, which is why Saga et al. designed a "daily support robot" [51] for use by children. The robot in this system is attached to a fixed track that is worn on the child's body. The robot uses accelerometer and gyroscope sensors, along with the known path of the track, to detect the child’s sitting posture. Saga et al. state that although their current system does not give any feedback, they aim in the future to have the robot help children improve their sitting posture by moving to the middle of their back and moving back and forth along the track there whenever the child slouches.

Moving slightly beyond sitting posture and into more general activity trackers, there are many commercial wearables that use sensors to detect a variety of health-related metrics. Products such as the FitBit Tracker [17] and Jawbone UP3 [28] use accelerometers to detect the number of steps taken, activity intensity, and how restless the user is while sleeping. The Jawbone UP3, which is worn on the wrist, also includes a heart rate monitor. Neither of these products claim to monitor posture, however a similar system, the Lumo Lift [33], does. The Lumo Lift is fixed to the user’s clothing, around the upper chest area, and monitors typical activity tracker metrics (steps taken and activity intensity) and the user’s standing and sitting posture. It targets upper back and shoulder posture while a similar system, the Lumo Back [32], targets lower back posture. The Lumo Back is worn on a belt around the lower back. Since both Lumo Lift and Lumo Back are commercial products there is limited information on how they function internally, but they both appear to be based on accelerometer data similar to the other activity trackers. Both systems are individually calibrated to the user and to the posture the user is attempting to maintain. As with other wearables, the Lumo Lift has a significant advantage over instrumented chairs and video-based posture detection systems because it can be worn all the time, not just at the computer workstation. It is also less disruptive and more discreet than the wearable systems seen in academic literature.

However, we were unable to discover any academic works verifying the accuracy of the Lumo products.

2.4.3 Instrumented Chairs

The most common method for detecting sitting posture in an office setting found in the literature is to augment the office chair with pressure or force sensors. Applying force sensors to the surface of a chair, or to a cushion that can be added to the chair, is relatively simple and low-cost when compared to the other methods of posture detection already described. Also, force sensors usually have a very low profile and are unlikely to disrupt the working environment. Often, they can be applied to existing chairs, requiring little to no change to the user's working environment. The biggest limitation of these systems is that each augmented chair is limited to a single workstation, unlike wearables which are tied to a single user. However, since the survey conducted by Griffiths and Saponas [22] indicated that most individuals have a primary chair in which they spend most of their time, we do not consider this limitation to be a serious disadvantage. As with most posture detection systems, augmented chairs must be calibrated to each user to achieve acceptable recognition rates, which limits such systems to offices in which each worker has their own personal chair. The works in the literature that use augmented chairs can be distinguished primarily by where the sensors are placed, how many sensors are used, and how the postures are classified.

2.4.3.1 Sensors in Cushions

Several of the works in the literature use augmented cushions that can be moved between chairs. These systems are generally simpler than the systems described in the next Section since they have fewer sensors and do not need to be calibrated to the individual user. The PosturePad of Epstein et al. [15] is a cushion augmented with sensors, however the paper is vague on technical details so it is unclear precisely how the cushion works and how it is used. The PosturePad can detect forward and backward lean only. Another augmented cushion is the CushionWare of Liang et al. [31], which is used as an input for a racing video game and an electrical wheelchair. The cushion detects whether the user is leaning in one of three directions (forward, left and right) or sitting upright, and the car in the game, or the physical wheelchair, moves in the same direction. Sitting upright causes the car or wheelchair to stop. Both of these systems demonstrate the most important limitation of augmented cushions: they can only be used to detect lean direction rather than the user's entire sitting posture. Another limitation is that it is difficult to ensure that the cushion is used properly by every user, since it is not fixed to a particular position on the chair.

2.4.3.2 Sensors Under Chair Base

When applying sensors to an actual office chair, there are two options for sensor placement: the surface of the cushions, addressed in the next Section, and the underside of the chair near the support. When four force sensors are applied under the base of an office chair, the centre of pressure (CoP) of the user sitting in the chair can be calculated. This allows for some posture classification, however spinal posture must be determined by a mathematical model instead of being detected directly. Two connected works out of the Upper Austria University of Applied Sciences implemented this method of posture detection, one by Haller et al. [24] and another by Schrempf et al. [52]. Haller et al. [24] used the CoP system to detect static sitting, which is when the user's sitting posture is unchanging for a period of time. They did this by calculating how much the CoP moved over time and considered any period of static sitting longer than 5 minutes as poor posture. Haller et al. also compared the three feedback modalities (Physical, Digital and Vibrotactile) in context of static sitting; this comparison is discussed further in Section 2.5.

Schrempf et al. [52], added a biomechanical model to this CoP method that scores the user's current sitting posture by the amount of biomechanical strain it creates on the user's body. Their system was not intended to classify the sitting posture into discrete categories, but rather to score the posture along a "healthiness" continuum. This is an unusual approach, and Schrempf et al. found their model could reliably estimate externally-measured parameters of sitting posture such as the force applied to the foot support. However, no other work was found that repeated this approach and so the validity of this model is less certain compared to the more prevalent methods of posture detection and classification used with sensors applied to the chair surface.

2.4.3.3 Sensors Applied to Chair Surface

The final method of posture detection with an augmented office chair found in the literature involves applying sensors to the surface of the chair. Several influential papers that are often cited in other posture detection literature fall into this category. The chairs in these works vary in the number of sensors used from a pair of sensors, one in the chair base and another in the chair back [10], to up to 19 sensors across the chair surface [37]. More recent papers have converged on 6-8 sensors [27, 62] as sufficient to detect sitting posture.

The earliest posture detection works, such as the work of Tan et al. [55], used sensor sheets laid over the office chair to detect sitting posture. One of the first works that applied individual sensors to particular locations on the office chair is that of Daian et al. [10], published in 2007. Their single force sensor in the chair base detected presence in the chair, and the single sensor in the chair back detected sitting posture. The data from these sensors was compared to thresholds that Daian et al. determined through repeated pilot testing. The user's posture was considered to be poor if these thresholds were exceeded for more than 20 seconds. Daian et al. also measured sitting times with their augmented chair, and included break reminders in their study that were triggered after 20 minutes of constant sitting. Any period of standing of over 5 seconds was considered a break. Daian et al. tested their chair with a limited number of participants and presented no statistical results regarding the accuracy of the sensors.

To establish the validity and accuracy of detecting posture using force sensors placed on the chair surface, Mutlu et al. [37] determined the near-optimal placement of 19 sensors based on posture data from a sensor sheet that covered the entire chair surface. Their aim was to create a classification system that could generalize between different users and therefore would not have to be calibrated to each user. This is a higher standard than aimed for in most subsequent works. In the process of determining near-optimal placement for their 19 sensors, Mutlu et al. found that the location of the sensors was more important to the accuracy of the posture classification than the intensity of pressure on any given sensor. They also found that there was more variation in posture data between users than between postures for a single user. Mutlu et al.'s system classified the user's posture into one of 10 possible postures, which they based on Tan et al.'s work [55]. Only one of the 10 postures was deemed to be a good posture.

Using a similar style of office chair, Zheng et al. [62] reduced the number of sensors on the chair surface from 19 to 7, with only 2 on the chair back and the rest on the chair seat. Rather than use a mathematical model, Zheng et al. chose the sensor locations based on the most distinguishable areas of the human body. This paper is one of the most cited by other works on posture detection, and it was also one of the first to move beyond posture detection and include a feedback component. Thus it is one of the first persuasive posture detection systems. Zheng et al. chose to use vibrotactile feedback, and this aspect is discussed more in Section 2.5.2. Similar to Mutlu et al. [37], Zheng et al. classified their participant's posture into one of 10 defined postures, of which only the "upright" posture was considered to be a good posture. Their participants sat in each of these 10 postures, and the sensor data was then recorded as the reference for that posture. After these references were set, Zheng et al. classified the participant's current posture by calculating the mean squared error between the sensor data for the current posture and that of each of the 10 possible reference postures. The reference posture with the lowest mean squared error was considered to be the matching posture only if it was below a certain threshold, otherwise the posture was classified as "other". They achieved an 86% accuracy rate for detection among the 10 different postures, and a 94% accuracy rate for detecting among only "upright", "slouching", "leaning forward" and "leaning back" postures. The latter accuracy rate excluded the side-to-side leaning and leg-crossing postures.

One of the only commercial posture-detecting chairs is the Axia SmartChair [6], which has a similar sensor setup to Zheng et al.'s chair [62], with four sensors in the chair seat and two in the chair back. The chair detects the users' posture once per second and classifies it as one of seven possible postures [21]. The SmartChair incorporates a small label on the side of the chair that shows a diagram of the chair with small LEDs that light up according to the user's dominant posture over the last hour. Each chair must be calibrated to the individual user's "pressure profile" [13] in a reference posture before first use. The SmartChair has been used in multiple studies, including the work of Netten et al. [38] where it was used to determine presence in chair when studying adherence to sitting guidelines.

There are several other augmented chairs in the literature that use sensors on the chair surface. The BITAIKA system developed by Ishimatsu and Ueoka [27] uses 6 sensors in the base of a chair to determine when the user's body is balanced unevenly on the chair. Ishimatsu and Ueoka ascribed all instances of this unevenness to the user crossing their legs. The spinal posture in this case was detected with the Kinect, as discussed in Section 2.4.1. Martins et al. [35] used inflatable pressure cells to detect posture instead of force sensors; the locations of these cells on the chair were influenced by Zheng et al.'s work [62]. Martins et al. classified their participants' postures into one of 11 pre-defined postures using Artificial Neural Networks and achieved recognition rates of over 70%. The HealthChair, developed by Griffiths and Saponas [22], does not detect sitting posture but does use force sensors on the chair back to detect respiratory rate. They conducted a survey of 550 people and found that 67% of them reported using the chair backrest. This supports placing sensors on the chair back as part of a posture detection system, as long as there are other sensors to account for those situations when the user is not in contact with the chair back.

2.5 Changing Posture with Feedback

The majority of papers use chair-based sensors to detect posture rather than any other method. However, the feedback modality used to communicate information about that posture to the user is much more varied. Examples of digital feedback, physical feedback and vibrotactile feedback are all common, both alone and in works that compare different feedback modalities. Only one work was found that did not use one of these three feedback modalities: Martins et al. [35] used inflatable pressure cells to detect the user's posture, and then changed the amount of air in the cells to cause the user slight discomfort and thus prompt them to change their posture. Their work focused mainly on posture classification and did not include tests of the feedback system. Another unusual feedback modality was proposed by Saga et al. [51], with their "daily support robot" moving at a point on the wearer's body to indicate an area of poor posture, however this feedback was entirely theoretical and was never implemented.

2.5.1 Physical Items

Several works used physical items on the desk to provide feedback to the user, often noting that this feedback modality is likely to be less disruptive to the user's workflow than digital feedback. Daian et al.'s early posture detection work [10] includes a physical agent on the desk that turns its back whenever the user adopts "inadequate posture" and faces forward when their posture is "adequate". The agent also moves back and forth to indicate that the user should exercise. It is accompanied by audible spoken messages to prompt breaks and stretching, to verbally reward good posture, and to give tips on how to improve bad posture. Daian et al. chose this feedback modality for the emotional component and because of the reduced disturbance to workflow compared to digital feedback. They found that participants spent more time in good posture and less time in bad posture when the feedback was enabled. However, due to the limited sample size no statistical tests were done on these results. Also promising are the results of their post-study survey, which found that participants showed little confusion and no irritation towards the physical agent.

Another work that includes affective, or emotional, feedback is the perFrames system of Obermair et al. [40]. The feedback in this system consists of a digital photograph of a loved one on the user's desk. The photograph changes from a smile to a frown depending on user's posture. It is a Wizard of Oz system so there is no posture detection component beyond an expert who monitored the participants through a webcam and triggered the changes in the photograph. Obermair et al. found that some of their participants disliked the lack of information provided by such simple positive/negative feedback, which is one of major drawbacks of physical feedback in general. However, the authors also found that most of their participants felt that the perFrame did not distract from their work and that the perFrame was aesthetically pleasing.

Haller et al. [24] compared all three of the main feedback modalities (physical, digital and vibrotactile) in an attempt to find the best way to interrupt users to improve their sitting posture. Each of the feedback modalities they compared will be discussed in the appropriate section in this work. Haller et al.'s physical feedback was a plant toy sitting on the user's desk that wilts as the user's posture deteriorates. It also shakes its leaves to trigger the user to exercise. This feedback modality was rated as the least disruptive of the three studied by the participants, although some participants took a while to notice the physical avatar had changed. Haller et al. also administered a modified

NASA-TLX [26] survey to measure the experience of the interruptions. While the feedback modality did not have a significant effect on most of the dimensions measured by the NASA-TLX, Haller et al. found that the physical feedback was significantly less disturbing to the participant's workflow than either digital or vibrotactile feedback. Importantly, Haller et al. also found that the type of task the participants were engaged in had a more significant effect on their likelihood to postpone an exercise session than did the feedback modality.

2.5.2 Vibrotactile Feedback

The work of Zheng et al. [62] has promoted vibrotactile feedback as an effective modality that is less disruptive than digital feedback. However, some works have also found vibrotactile feedback to be irritating to users. As with physical feedback, there are limits on the amount of information that can be conveyed to the user through vibrations. Zheng et al. [62] placed vibrotactile actuators on 4 locations on the chair back and 2 on the chair base to communicate to the user precisely how they should adjust their posture. Each actuator was associated with a particular aspect of poor posture, e.g. slouching, crossing the left or right leg, or leaning to the left or right. The vibration of each actuator was intended to indicate a particular posture correction to the user, although these were not explained explicitly to the participants. Zheng et al. incorporated a time delay to prevent the actuators from vibrating constantly. They found that their participants were able to quickly map the vibration locations to the appropriate action that would cause the vibration to cease. They also found that the cycling through sessions of 5 minutes of feedback alternatively enabled and disabled was sufficient to change the participants' posture significantly compared to when there was no feedback at all. Despite these frequent, as-needed reminders to maintain a particular, upright posture, participants would still lapse into slouching and other poor postures in between vibrations. Zheng et al.'s results suggests that although intermittent, as-needed vibrotactile feedback is sufficient to effect short-term change, the change does not continue for long once the feedback is removed. Whether this is the case over longer-term exposure to regular feedback is unclear.

A simpler system that also uses vibrotactile feedback is the PosturePad of Epstein et al. [15]. This system is an instrumented cushion that detects only whether the user leans too far forward or back, at which point it vibrates and emits an audible beep. The PosturePad feedback is continuous, with no time delay, and always enabled. In a study with adolescents, Epstein et al. compared this continuous reminder with verbal reminders from a teacher at regular intervals. The participants' posture was assessed visually by parent volunteers. Their results indicate that although participants' posture improved immediately after the verbal reminder, it deteriorated quickly, with no significant difference between the initial posture and the posture 16 minutes after the verbal reminder. However, the continuous, contextual feedback provided by the PosturePad caused students to spend significantly more time in a good posture over a 10 minute period. Although this is a relatively simple measure of posture and a very short-term study, it does suggest that feedback based on the user's current posture is more effective than periodic reminders absent any context.

Haller et al.'s comparison of all three feedback modalities [24] included a vibrotactile component however there were few technical details in that paper as to where the vibration occurred. The authors did cite the Zheng et al. study [62], which suggests the vibration occurred on the chair. Four participants in Haller et al.'s study found continuous vibrations extremely annoying and mentioned they might shut off the vibrations with longer-term use. This could explain why participants reacted more quickly to the vibrotactile feedback over the other two modalities. The results from the NASA-TLX survey suggest that the vibration feedback was more disturbing to participants' workflow than digital or physical feedback. This contradicts Zheng et al.'s [62] finding as most of their participants reacted favourably to the vibrotactile feedback. The contradiction could be due to the differences in the length of the vibrations, since Zheng et al. cycled between short periods of enabled and disabled feedback while Haller et al.'s feedback appears to have been always-enabled. The cycles would create periods of rest that could have improved participants' feelings toward the feedback. Zheng et al. did not measure any more objective metrics of workflow disturbance, so they cannot be compared with that aspect of Haller et al.'s results.

The commercial posture-detecting chair described in Section 2.4.3.3, the Axia Smart Chair [6], also uses vibrotactile feedback. The work of van der Doelen et al. [13] mentioned that this vibration occurs under the upper right leg of the user near the front of the chair base, and that the vibration consists of four short pulses over 4 seconds, emitted at most once every hour. This is a much milder vibrotactile signal than the three studies already discussed in this Section. Van der Doelen et al. studied the effectiveness of the SmartChair feedback [13]; they looked at data from 21 participants over a six-week period and found that they spent significantly more time in the reference posture in the 10 minutes after the vibration signal than in the 10 minutes before. They also observed several situations in which the participants got out of the chair immediately following a vibration signal. This was an unintended effect of the vibration but van der Doelen et al. noted that it could be considered healthy behaviour as it was interrupting periods of extended sitting.

The SmartChair was also used in a study by Goossens et al. [21] over a 4-week period. They compared participants in three groups: those who received no feedback (the control group), those who received instructions on the optimal use of the SmartChair, and those who received instructions and had the feedback features of the SmartChair switched on. In the first week after the instructions were delivered and the feedback was switched on, participants in the two experimental groups spent significantly more time in the reference (i.e. good) posture than the control group. Over two more weeks this effect deteriorated, although the SmartChair feedback group still maintained slightly higher amounts of time in the reference posture than either the instruction-only group or the control group. These results suggest that vibrotactile feedback may lose its effectiveness over time as the users adjust to the vibrations, but that it is still an improvement over one-time verbal instructions.

Two more commercial products that use vibrotactile feedback are the Lumo Lift [33] and the Lumo Back [32], described in detail in Section 2.4.2. When either product is put into “coaching mode” it will vibrate whenever the user adopts a poor sitting or standing posture. The reference posture is set by the user and poor posture is defined as any posture that does not match this reference. This feedback can be turned off at any time by the user but is otherwise continuous. There may be a time delay to the vibration but that is not clear from the information provided by the manufacturers. No academic literature was found that examined the effectiveness of this feedback.

2.5.3 Digital Feedback

In general digital feedback includes any feedback that appears on a screen, whether it is a computer monitor, a standalone screen, or a smartphone. Several of the posture-detection systems found in the literature place their digital feedback on a traditional computer monitor. These types of digital feedback can include anything from small taskbar notifications to windows of varying sizes. Digital feedback is thought to be the most disruptive of the three feedback modalities since it is a purely visual medium, and thus uses the same cognitive resources as most of the work that computer users are engaged in. This may explain why many authors of posture detection systems have chosen vibrotactile feedback over digital feedback. However, digital feedback uses communication methods that are already familiar to office workers, such as pop-ups and taskbar notifications. This makes the feedback easier for users to understand and react to.

Two systems in the literature fit only the more general digital feedback definition since they are not implemented on computer monitors. First, in the Augmented Reality Posture Mirror (ARM) of Taylor et al. [56], the digital feedback appears over top of an image of the user that is displayed on a vertical screen, mimicking the effect of a mirror. The feedback appears as either a mist of increasing intensity, depending on the degree of deviation from the user's reference posture, or as points of colour that indicate the exact area of the poor posture. In the latter case the image of the user gives context to the digital feedback. The ARM system was designed to work with both sitting and standing posture, which may be one reason Taylor et al. chose to work with an independent screen. However, the technique could easily be applied to a pop-up window on a computer monitor as well.

The other system that does not use a computer workstation monitor is the Lumo Lift [33]. This device provides digital feedback through a smartphone app, in addition to the vibrotactile feedback on the wearable itself. This more portable form of digital feedback is appropriate for a system designed to be worn constantly and not only at a computer workstation. The Lumo Lift app provides mostly general feedback on the user's posture and activity level in the form of 3 categories: Remarkable, Good or Slouchy for posture, and Super Active, Active or At Rest for activity. Users can see more information on how many hours they spent in good posture and how many steps they have taken, shown next to their daily goals. The app does not include any truly specific feedback that indicates what the user's posture actually looks like compared to the reference.

One system that does use a computer monitor to communicate digital feedback is the BITAIKA system of Ishimatsu and Ueoka [27]. This combination of Kinect and augmented chair seat system uses a pop-up on the computer screen to display the side view of the user, taken from the Kinect camera. The line of the user's centre of gravity is highlighted in red, while the ideal centre of gravity line is highlighted in green. The user in this system is expected to move until the red line is on top of the green line; if the user achieves this for 2 seconds the pop-up minimizes itself. The authors found that their participants spent significantly more time in good posture with this feedback than without it, which is a promising result. However, the participants also expressed irritation with the frequency of pop-ups (an average of 16 times in 30 minutes) and complained of back pain and fatigue. The authors dismissed the back pain as evidence that the participants were activating unused muscles, but any system that causes pain or fatigue over such a short period would be likely to cause severe problems with longer term use. The feedback in this system seems to be too frequent and too coercive for the participants' comfort, since the pop-up can only be closed by matching the ideal centre of gravity line. Whether a less frequent or less coercive feedback method would have been equally effective at changing participants' posture is still an open question.

Haller et al.'s comparison of all three feedback modalities [24] included a digital component that combines a taskbar icon and a pop-up window. The taskbar icon mimics the physical avatar on the desk; the plant image "wilts" as the participant's period of static sitting lengthens. The pop-up window pops up on the right side of the screen and resizes the main window the user is interacting with. Its appearance is meant to prompt the user to perform an exercise, and if the user agrees then the main window shrinks to the left and the exercise window expands to fill the majority of the screen, displaying an exercise video for the user to follow. Once the exercise is complete the exercise window disappears. Participants in this study took significantly longer to start an exercise session and significantly longer to return to their main task after the exercise when they received the prompt to exercise from digital feedback compared to vibrotactile or physical feedback. Haller et al. also found that the participants often looked at the taskbar icon when switching between applications, suggesting that it was less of a disruption to their workflow. However, the NASA-TLX results showed that the digital feedback was more disruptive than the physical feedback and less disruptive than the vibrotactile feedback. From these results digital feedback can be considered a "middle ground" of feedback modalities. Haller et al. did not present any data on whether any of the feedback modalities they compared were actually effective in changing their participants' posture.

The only commercial posture-detecting solution found using digital feedback on a computer monitor is the Philips ErgoSense Monitor [47]. It is not clear from the documentation how the digital feedback in this system works, but it appears to be small on-screen pop-ups originating from the taskbar. Taskbar notifications are considerably less disruptive than pop-up windows, and this choice of a low-disruption feedback is understandable in a commercial product that does not want to irritate its users. No academic literature regarding this system was found, so there is no evidence as to whether or not the feedback in this system is effective at changing users' posture.

Chapter 3 The PostureChair System

To address our research questions, a system was designed and built that detects the user's posture in an augmented office chair, and displays digital posture feedback on the user's computer screen. The office chair and associated software together make up the PostureChair system. Two instances of this system were built to run two participants simultaneously, and to have a backup in the event of equipment failure. The following Sections give an outline of the design process for the PostureChair system, in the context of the field of persuasive technology and previous academic works in posture detection and feedback. The hardware and software components of the system are then described in more detail.

3.1 Designing Persuasive Technology

Designing successful persuasive technology can be difficult since users do not always respond as expected to attempts at persuasion. To aid in the development of the PostureChair system we followed the first 6 steps of BJ Fogg's design process for persuasive technology [18], as outlined below:

1. *"Choose a simple behaviour to target."* Posture detection systems can also be used to measure sitting times and frequency of sitting breaks by detecting presence in the office chair. Interrupting long periods of sitting with standing breaks is likely as important in terms of overall health as improving posture. However, a system that attempts to target both sitting posture and break times would be less effective than a system that targets a single behaviour. The PostureChair system explicitly targets only the improvement of posture and targeting sitting breaks is left as possible future work.

2. *"Choose a receptive audience".* In testing the effectiveness of the PostureChair system we sought out participants among the students and employees of the local university. These participants were selected partly for convenience, but also because they are generally moderately heavy computer users and thus could be receptive to technology aimed at improving their posture. In phrasing the call for participants in terms of improving posture and healthy workplace behaviour, we were more likely to recruit participants that are interested in these topics. While this may cause a slight bias in our results in favour of changes in posture, we feel this is acceptable since it reflects the probable users of such a system in a real-world application of this work.
3. *"Find what is preventing the target behaviour: lack of motivation, lack of ability or lack of a well-timed trigger".* We believe that improving sitting posture is mostly a case of a lack of a well-timed trigger, or reminder. Information on sitting posture has been well-publicized in government and corporate literature for several years, and adjustable chairs and desks are becoming common, so lack of ability is unlikely to be the main issue. Based on our anecdotal conversations with other office workers most people are aware of ergonomics and the consequences of poor posture, and many actively try to sit ergonomically, making it unlikely that there is a lack of motivation. The PostureChair system was designed to provide a well-timed trigger to allow users to improve their posture while maintaining focus on their work.
4. *"Choose a familiar technology channel".* As shown in Section 2.5, there is no one method of feedback that is considered ideal for persuasive posture-detecting

systems. One of the rationales for choosing digital, on-screen feedback for the PostureChair system was that it would be more familiar to our participants compared to vibrations in the chair or physical items on the desk.

5. *"Find relevant examples of persuasive technology".* This step is an obvious one for an academic work, and we have previously described the results of our search for relevant examples in Sections 2.4 and 2.5.
6. *"Imitate successful examples".* Throughout the rest of this Chapter we note where we took inspiration from relevant works on posture-detecting chairs and persuasive digital feedback.

The final two steps in BJ Fogg's design process relate more to the actual development process (7. *"Test and iterate quickly"* and 8. *"Expand on success"*) and so are more relevant to a commercial development process than to this academic work.

3.2 PostureChair Design Overview

The PostureChair system performs both posture detection and posture modification through persuasive feedback. It is made up of two components: an augmented office chair and Windows-based software. From our review of the literature, covered in the previous Chapter, it is clear that various posture detection methods have already been thoroughly studied and tested for validity. As a result the posture detection aspect of this system was designed based on the most influential works in the literature, while the posture modification aspect was designed independently and was the focus of the testing.

3.2.1 Posture Detection

The PostureChair system detects posture using an office chair augmented with flat force-sensitive resistors. These sensors were chosen as they appeared to be the most reliable and least disruptive method for a computer workstation context. The layout of the sensors on the office chair was influenced by the work of Zheng et al. [62] and the Axia SmartChair [21]. As with most previous works, the PostureChair requires calibration to individual users and to ergonomic reference postures. Only Mutlu et al.'s work [37] attempted unknown user classification, and we consider this to be an unnecessary distraction for a work that focuses mainly on the effect of the feedback, rather than the power of the classification algorithm.

The data from the force sensors is communicated to the computer software at a frequency of 1 Hz. The software then classifies the user's current posture as either "good" or "bad", relative to the reference posture. More details on this algorithm can be found in Section 3.4. Although most of the augmented office chair works we reviewed classified the user's posture as one of 7-10 possible postures [13, 21, 35, 37, 55, 62], only one of those postures (the "upright" posture) was ever considered to be an ergonomic, or "good" posture. Since we are more concerned with the improvement of posture in our system, we have restricted our classification of the user's posture to a binary classification system of "good" and "bad". We consider a "good" posture to be one that matches the reference posture. The reference posture is recorded at the beginning of every study and is unique to each participant. The participant was instructed to follow the recommended guidelines for an ergonomic sitting posture, as outlined in Section 2.3.1, as closely as possible given the limitations of the office chairs and desk setup. Although we followed the ergonomic guidelines, the reference posture can actually be set to any sitting posture the user chooses. Accordingly, this study focuses on whether the feedback given by the PostureChair system can successfully alter the participants' posture to match their particular reference posture.

3.2.2 Feedback

As discussed in Section 2.5, the different modalities of feedback at a computer workstation have been well studied but the effectiveness of the feedback has not been demonstrated to the same extent. Daian et al. [10] examined the effectiveness of a physical avatar on the desk and Zheng et al. [62] did the same for the vibrotactile feedback they employed. The Axia SmartChair has been shown to be effective in the 10 minute period after a vibration signal is emitted [13], however over the longer term (2-3 weeks) this effect seems to deteriorate [21]. The BITAIKA [27] system includes digital feedback that requires the user to match the ideal posture before the window could be dismissed and they could return to their work, which in our opinion is more coercive than persuasive. These are the only works found that examined the effectiveness of any kind of feedback on sitting posture and none of them use persuasive digital feedback. The PostureChair system seeks to address this gap and determine if persuasive digital feedback can be effective at modifying computer users' sitting posture.

Digital feedback in the PostureChair system was chosen for several reasons. First, we believe it is the most appropriate feedback method for a computer user, since it requires little disruption to the workstation setup and uses metaphors and communication methods that the user is likely already familiar with. Second, Haller et al.'s comparison of the three main feedback modalities (digital, physical and vibrotactile) [24] showed digital feedback to be less disruptive than vibrotactile feedback. Haptic feedback such as vibrotactile has been the most commonly-chosen feedback modality in many other posture detection systems. Finally, digital feedback was selected precisely because it has been otherwise neglected in the literature.

Epstein et al. [15] showed that contextual feedback is far more effective than timed reminders, although they looked at vibrotactile rather than digital feedback. As a result, the feedback pop-ups in the PostureChair system appear only when the time and severity thresholds for bad posture have been exceeded. These feedback windows take up most of the centre of the screen and cannot be minimized, so the user cannot ignore them. When the user dismisses the feedback window the thresholds are reset. The user is not forced to change their posture in order to dismiss the window, which keeps this a persuasive, rather than coercive, system. In a real-world application the different thresholds for bad posture and the amount of information presented in a feedback window could be customized by the user; for the purposes of comparison these aspects were kept constant throughout the study.

Our second research question examines how much information needs to be present in digital feedback for it to be effective. To address this question the PostureChair system has two types of feedback: Generic and Specific. The Generic feedback is a simple indication that the user's posture has deteriorated, while the Specific feedback gives more information to the user that indicates how they can improve their posture. Digital feedback allows more complex information to be communicated than either physical or vibrotactile feedback, which can be considered an advantage of this feedback modality. The two feedback types in the PostureChair system were inspired by the two types of feedback in the Augmented Reality Mirror of Taylor et al. [56], and were included in our system to determine if this additional information is necessary to modify posture and whether it is helpful to users in any way.

3.2.1 Persuasion

Section 2.1 described how persuasive technologies can be categorized by their method of persuasion or by the type of change and the outcome they are attempting to achieve. The PostureChair system falls under four of the "method of persuasion" categories described by BJ Fogg [19]: Tailoring, Suggesting, Self-Monitoring and Surveillance. The PostureChair system *tailors* information to the context by presenting feedback on the user's posture only when poor posture is detected. The Specific feedback window *suggests* better posture by identifying precisely where the poor posture has been detected. The same feedback window also updates in real-time, giving the user an opportunity to *self-monitor* and adjust towards the ideal posture goal. Both feedback windows, along with how visually different the augmented office chair is from a regular office chair, alert the user to the *surveillance* that the chair is carrying out. All four methods of persuasion rely heavily on feedback.

In terms of Oinas-Kukkonen's model of persuasive technologies [41], with the PostureChair system we are attempting to achieve compliance (C-Change) in the short-term, with the goal of altering ingrained sitting posture behaviour (B-Change) over the long-term. Attempting to alter sitting posture habits, which have often developed over years, classifies the PostureChair system as an A-Outcome system.

3.3 Hardware: The Augmented Office Chair

The augmented office chair is composed of a typical office chair, with 8 force-resistive sensors (FSRs) on its surface, connected to an Arduino Mega 2560 [3]. The office chair was chosen partly for convenience as it was readily available, but also because it has some degree of adjustability. The armrests of the chair are vertically adjustable, the lumbar support has one axis of adjustment, and the seat back reclines on a single-point pivot. This last feature changes the tilt of the entire chair, including the seat, relative to the floor. The seat height relative to the floor is also adjustable. The chair conforms to the "minimum critical chair features" list in the Haworth Ergonomic Seating Guide [57], except that it does not have any seat depth adjustment capabilities. Ideally, computer workers would have access to much more adjustable chairs than the one used as part of the PostureChair system. However, in reality, this chair reflects the type of chair that many workers actually use. During the user study, described in Chapter 4, the chair tilt was set to be parallel to the floor and the lumbar curve of the chair was set in the middle of the available range. The participants were able to adjust the height of the chair seat and the height of the armrests.

The Interlink FSRs [20] used are 1.75 inches square and are flush to the surface of the chair. FSRs are made up of two conductive traces that serve as electrodes and a resistive polymer that is attached on top with a spacer to prevent immediate contact with the electrodes. When force is applied to the FSR, the polymer makes contact with the electrodes and completes the circuit. More force compresses the polymer, increases contact surface area within the polymer and thus creates more conduction paths, decreasing the resistance. For these particular FSRs the resistance is infinite when no force is applied and like all FSRs the resistance decreases as the force applied increases. The FSRs were fixed to the chair with tape and wired to the Arduino Mega, which was affixed to the back of the chair and protected from damage by a cardboard shield. The chair was then covered with a custom-sewn slipcover made of thin cotton fabric, to protect the sensors from wear without interfering with their ability to detect pressure. The front of the chair is shown in Figure 1 and the back of the chair is shown in Figure 2.

Figure 1. The augmented office chair, front view, without slipcover (left) and with slipcover (right)

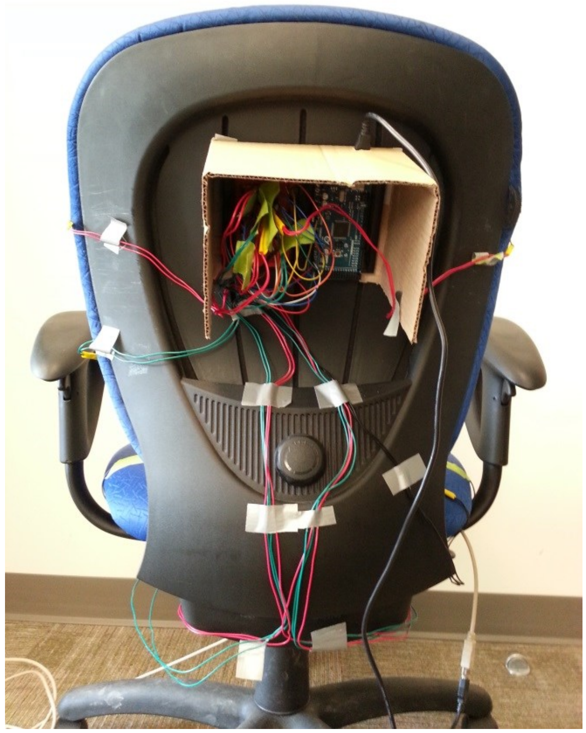

Figure 2. The augmented office chair, back view

The Arduino Mega is connected at all times to a Windows 7 desktop computer by a USB cable, which also provides power to the Arduino and FSRs. The FSR placements are partly based on Zheng et al.'s setup [62]: one FSR is placed under each sit bone, one under each thigh close to the knee, one at the apex of the lumbar curve and one behind each shoulder blade. The eighth FSR is placed at the centre rear of the chair seat to detect whether the user is sitting all the way back in the chair. The sit bones are the two bony swellings upon which most of a person's weight rests when sitting. The exact locations of the FSRs on the chair are based on the body and limb measurements of the 50th percentile of US Adults [12]. Only male measurements were available for the locations sit bones, so the distance between those two sensors, shown 18cm from the chair back in Figure 4, is based on male measurements only. All other sensor locations are based on the average of both male and female measurements. The exact positions of each sensor can be found in Figure 3 and Figure 4.

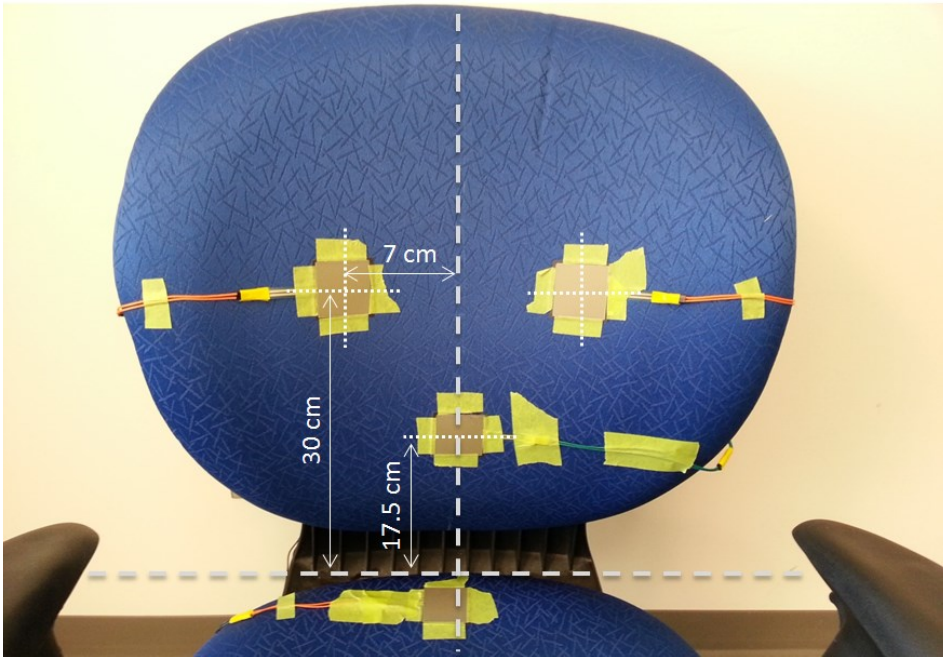

Figure 3. Positions of the FSRs on the chair back. The strong dashed lines represent the centre line of the chair back and the rear of the chair seat. The light dotted lines represent the centre points of each sensor.

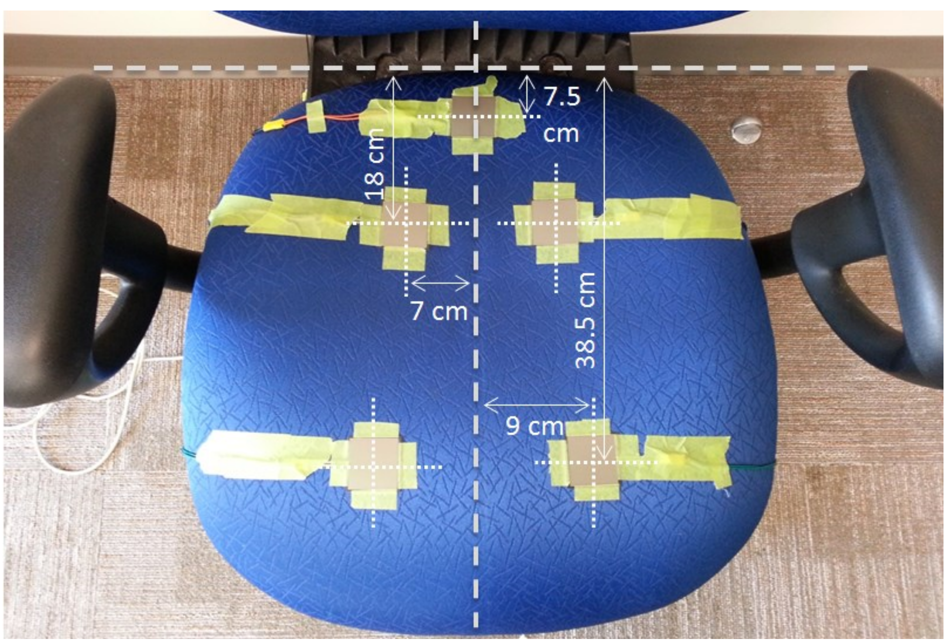

Figure 4. Positions of the FSRs on the chair seat. The strong dashed lines represent the centre line of the chair back and the rear of the chair seat. The light dotted lines represent the centre points of each sensor.

3.4 Software: Classification and Feedback

All of the posture classification and feedback to the user occurs in the PostureChair software running on the desktop computer. The PostureChair system also includes software on the Arduino Mega, which is written in the Arduino programming language [4]. The sole purpose of the Arduino software is to read the output of each of the 8 FSRs on the augmented office chair and communicate these readings to the computer software. The Arduino software combines the readings from each FSR into a single comma-separated line of text, which it then sends to the computer software once per second through the serial-to-USB adaptor built into the Arduino.

The PostureChair desktop software is a Windows Forms application, written in C# and developed in Visual Studio 2010. This software has two parallel threads: one that reads in the sensor data coming from the Arduino software and another that parses that data, classifies the posture and presents feedback to the user when necessary. Each time the software is run the user must calibrate the system by setting the reference posture in the Setup Window, shown in Figure 5. The user sits in the chair in a relaxed, neutral position, as far back in the seat as possible, with hips and knees flexed to 90°, feet flat on the floor and looking straight ahead, as specified in the posture guidelines described in Section 2.3.1 The user then clicks the "Record Baseline" button to record the current sensor data as the reference posture.

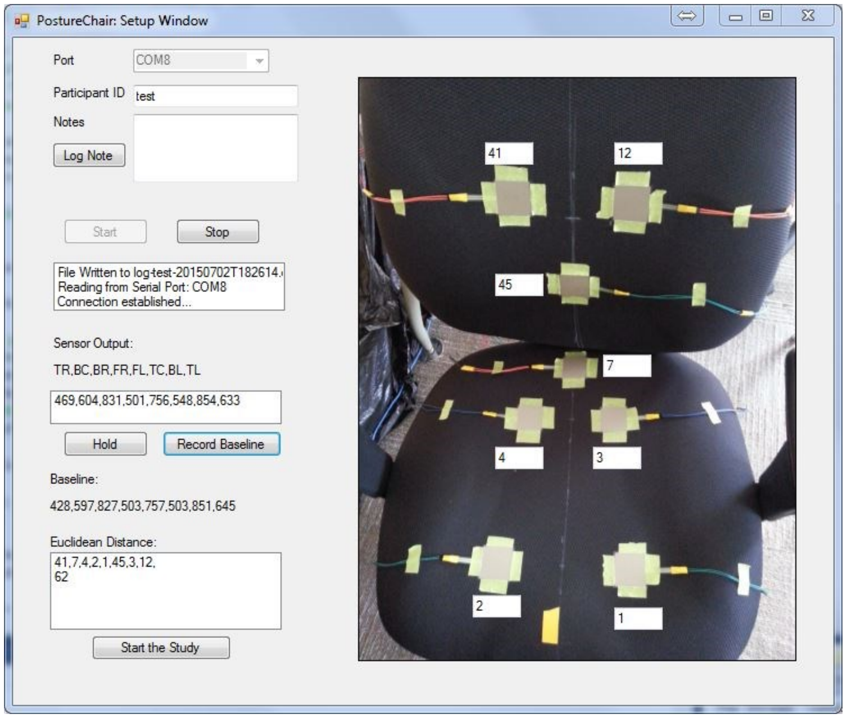

Figure 5. PostureChair Setup Window, shown after the reference posture (called the "Baseline" in this window) has been set.

The distance between the user's current posture and the reference posture is calculated using Euclidean distances. To achieve this, we consider each instance of sensor input as a point in Euclidean space. If all 8 sensors are considered together, the space has 8 dimensions. If a particular sensor is considered individually, the space has only 1 dimension. In either case, a Euclidean distance can be calculated between the reference posture, point x, and the current posture, point y. In 1-dimensional space, the Euclidean distance d(x, y) is the absolute difference between the two values x and y:

$$d(x, y) = |x - y|$$

In 8-dimensional space, if we consider x_1 to x_8 to be the 8 coordinates that describe the point x, and similarly for point y, then Euclidean distance d(x, y) can be described as:

$$d(x, y) = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2 + \cdots + (x_7 - y_7)^2 + (x_8 - y_8)^2}$$

This is more formally known as the L^2 norm. The Euclidean distance calculation is a non-weighted calculation, so all 8 sensors are considered to be interchangeable in the posture classification algorithm. In reality, the sensors on the chair seat usually experience a greater force than those on the chair back when a user is seated in the chair. It is also possible that the sensors may differ in how much their signals vary between good and bad posture. However, no works were found in the literature that would give a basis for weighting one sensor over another, so this simpler, non-weighted classification was chosen instead. Examining different sensor weightings is left for future work.

The software calculates the Euclidean distance between the reference posture and the current posture, for each sensor individually and for the 8-dimensional system overall, once per second. The Euclidean distance for the 8-dimensional system is then compared to the "posture threshold". If the distance is found to be less than this threshold, then the software classifies the posture as "good"; otherwise it classifies the posture as "bad". The software maintains counters to keep track of how long the user has been sitting in either good or bad posture. If the bad posture counter exceeds a time threshold of 60 seconds, and the feedback is currently enabled, the user is alerted to their poor posture with a feedback window. If at any point the good posture counter exceeds a time threshold of 15 seconds, then both counters are reset. This good posture time threshold allows the user to alternate very short periods of good and bad posture without triggering an excessive number of feedback pop-ups. A flowchart representing this classification algorithm is shown in Figure 6.

The posture threshold and time thresholds were selected after conducting a pilot study with the system. The goal of the pilot study was to select thresholds that generated a reasonable number of pop-ups so that participants in the main user study would have a chance to be affected by the feedback and improve their posture. The posture threshold had to be low enough to cause most participants to see at least one feedback pop-up, but not so low that they would be unable to prevent pop-ups from appearing by improving their posture. The pilot study was conducted with 3 participants who were asked to sit in the chair for 15-20 minutes while they completed their own work on a computer. They were instructed to sit in their normal posture and to try to forget they were being monitored. Their posture was recorded throughout this period and no feedback was given. Various posture and time thresholds were proposed based on a visual examination of the resulting posture data. The number of feedback pop-ups that would have appeared to each of the pilot study participants with each proposed posture and time threshold combination was then calculated. The results were compared to find a posture and time threshold combination that produced a reasonable number of feedback pop-ups for all 3 pilot study participants. These “middle-ground” thresholds were selected as the final posture and time thresholds used in the PostureChair system throughout the main user study.

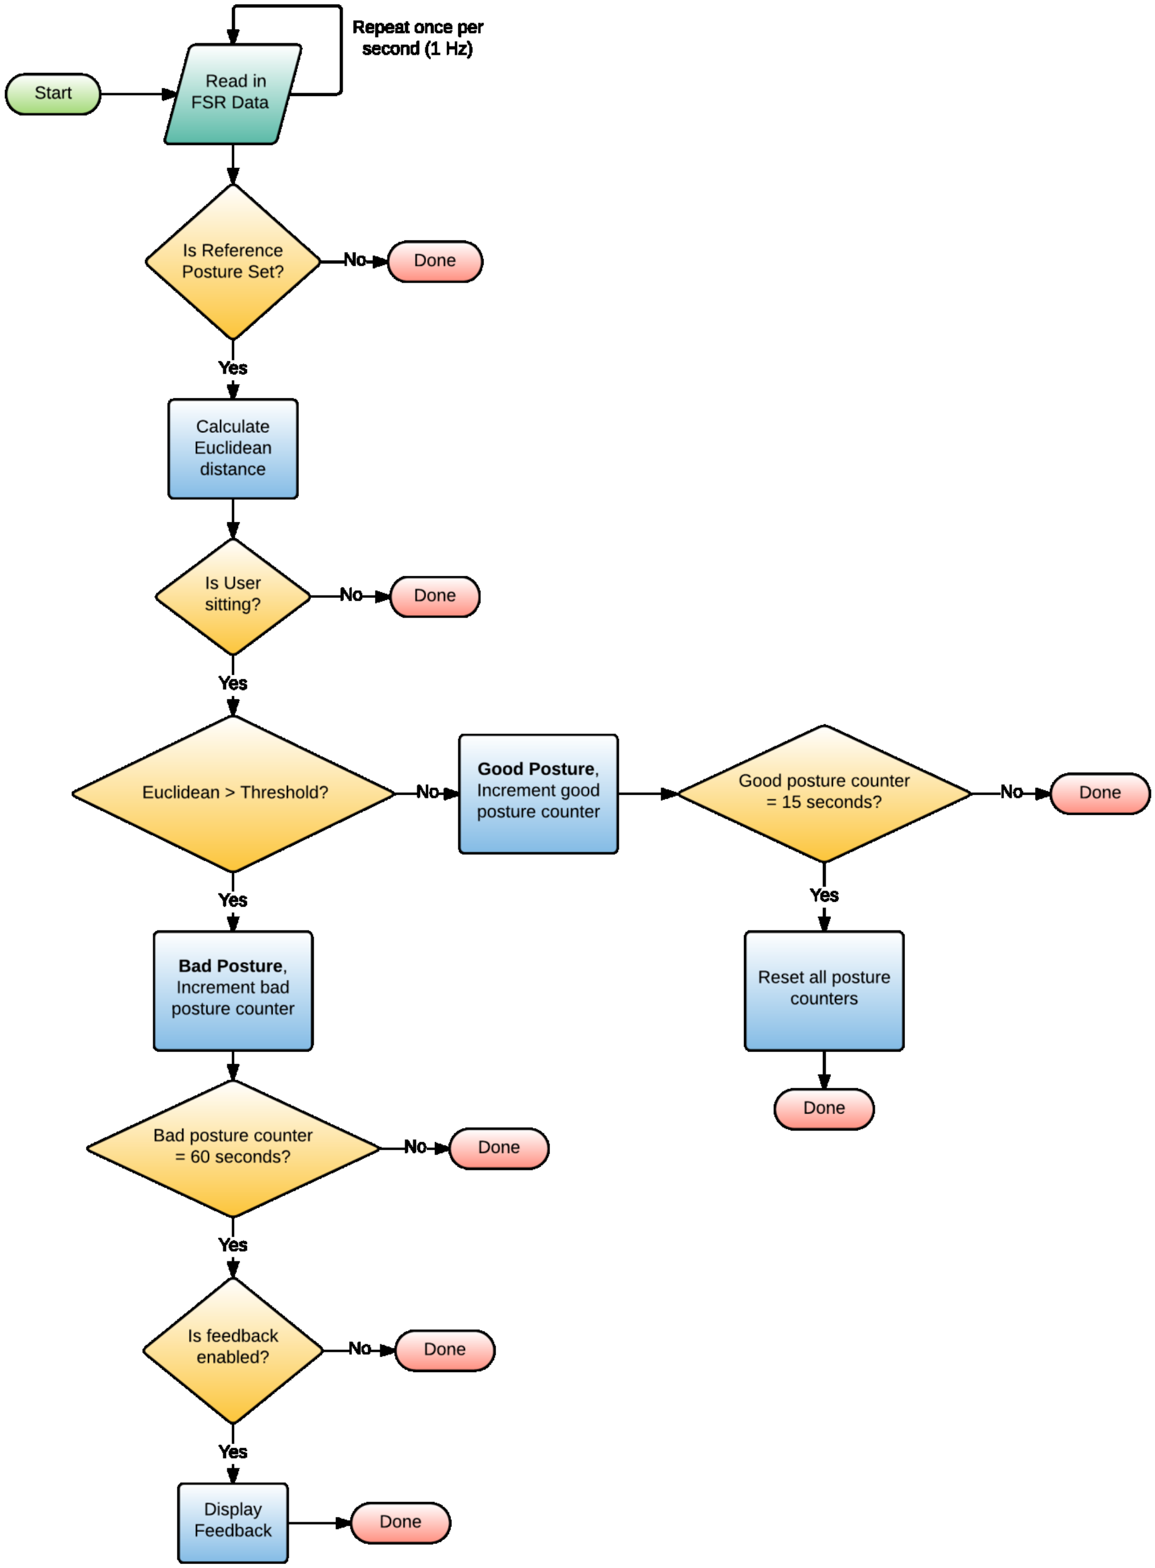

Figure 6. Flowchart representing the PostureChair posture classification algorithm

3.4.1 Feedback Windows

The PostureChair system has two different digital feedback windows that can appear when the feedback threshold is exceeded: a Generic window shown in Figure 7 and a Specific window shown in Figure 8. The windows have identical functionality and differ only in how much information is presented to the user. The feedback window appears in front of any other window the user has opened, and cannot be minimized or closed except through the "Ok, I got it!" button at the bottom of the window. This always-in-front aspect ensures that users will not miss the feedback if they are rapidly switching between tasks. The feedback window was designed to be seen as quickly as possible and to prevent users from ignoring the window, while remaining persuasive rather than coercive. It was not designed to minimize disruption to the user's workflow. When the window appears, the "Ok, I got it!" button is initially greyed out. The software records the user's posture at the moment the feedback window appears and activates this button only once the user has moved enough to cross a "posture change" threshold. This is a very low threshold and can be easily accomplished by the user. This ensures that the user cannot accidentally dismiss the window with the space bar or enter key while typing, since the feedback window appears without warning. With these features we can be sure that the user has always had a chance to see and register the information provided by the feedback windows.

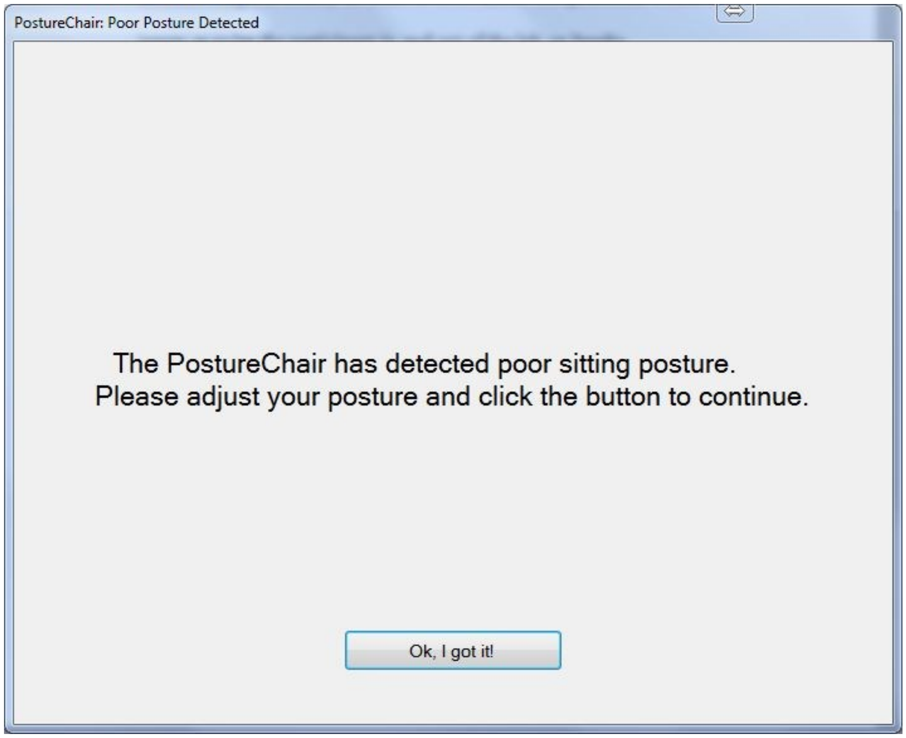

Figure 7. Generic feedback window

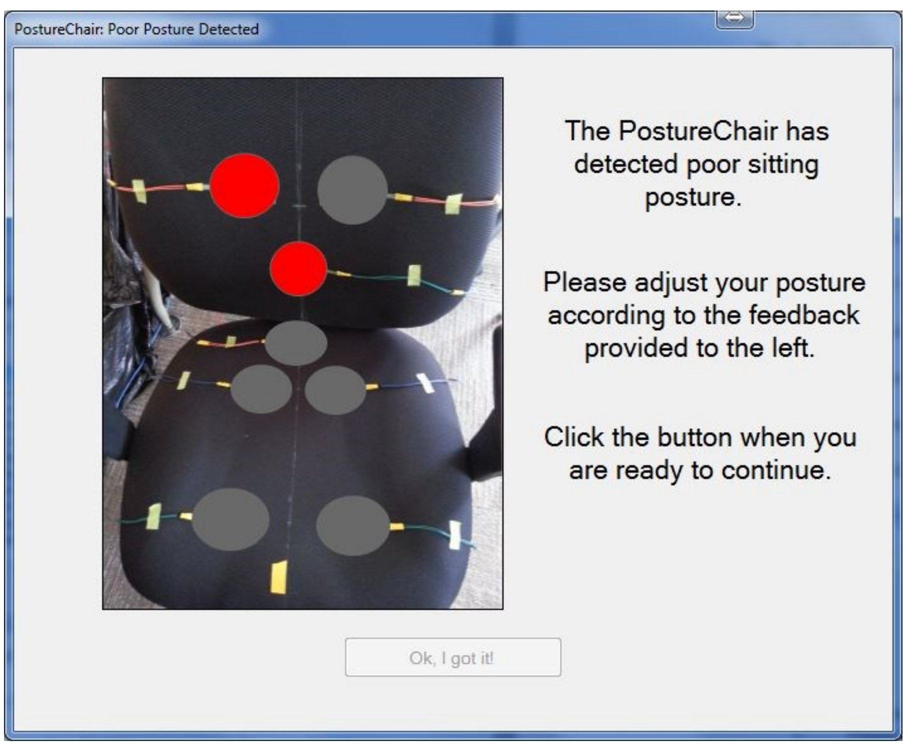

Figure 8. Specific feedback window

The Generic feedback window (Figure 7) presents only a simple message to the user that they have adopted a poor sitting posture. No information is given on how the user should adjust their posture.

The Specific feedback window (Figure 8) presents the same message to the user and additionally shows a picture of the augmented office chair with a grey circle representing each sensor on the chair. These circles turn red if the Euclidean distance between the current and reference postures for that particular sensor has crossed the single sensor posture threshold. The single sensor posture threshold was selected by the same process as the posture threshold for the 8-sensor system, as described above, except that the goal was to choose a threshold that would display the correct colour when the user was in a good or bad posture. The circle colours update continuously with each sensor reading (i.e. once per second) while the pop-up is open, so the user can adjust their posture and see how it affects individual sensor results. We hoped that this would allow users to experiment with their posture and determine which aspects of their posture caused the feedback window to appear when it did.

The only difference between the two feedback windows is this continuously-updating graphical representation of the sensor data. This allows us to compare how these two feedback windows affect the user's experience and actual posture, and to determine the effectiveness of the additional information provided in the Specific feedback window. The within-subject user study, as described in the next Chapter, was conducted to investigate these questions.

Chapter 4 Methods

The objective of our user study was to examine whether digital feedback was effective at altering our participants' sitting posture, and whether the extra information in the Specific feedback window was more effective than the simple message provided by the Generic feedback. We recorded objective measures, such as sensor data and the number of times a feedback pop-up window appeared, in logs generated by the PostureChair system. We also examined the users' experience with the system with a questionnaire that participants completed at the end of the study. The study was divided into three sessions of equal length. Feedback was disabled in the Baseline session to record the participant's unaltered posture. The Generic feedback was enabled in the second session, and the Specific feedback was enabled in the third session. This ordering was not counterbalanced between participants. Counterbalancing is generally considered to be good practice in user studies, and by not counterbalancing we risk having the learning effect of the Generic feedback influence the results related to the Specific feedback. However after careful consideration we believe that the learning effect of the Specific feedback far outweighs that of the Generic feedback and that by counterbalancing we would be more likely to influence the Generic feedback results than we would influence the Specific feedback results by not counterbalancing.

4.1 User Study

The user study was conducted on two Windows 7 desktop computers in a student computer lab at Carleton University. Participants were invited to come to the lab at a time of their choosing, and were encouraged to bring their own work to perform on the computer during the study. The setup of the study is shown in Figure 9.

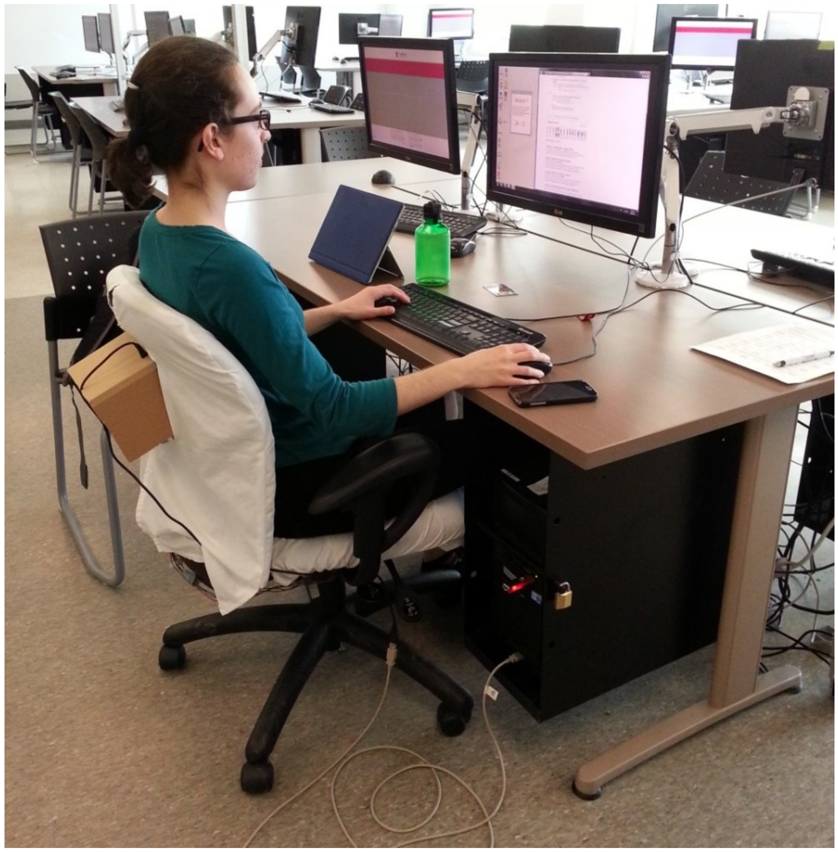

Figure 9. Participant in PostureChair during the Baseline session

4.1.1 User Study Outline

The following is an outline of a typical study:

1. The participant arrives at the computer lab and is briefed on the user study procedure. Photos of both feedback pop-ups are shown to the participant and their functionality is explained.
2. The participant signs a consent form (Appendix C) and takes a seat in the chair. The researcher then starts the PostureChair software and opens the Setup Window (Figure 5).
3. The participant is allowed to adjust the height of the chair to their comfort. The researcher adjusts the height of the monitor so that the top of the screen is in line with the participant’s eyes.

4. The researcher instructs the participant to sit as far back as possible in the instrumented office chair, in a relaxed fashion, with hips and knees flexed to 90°, feet flat on the floor and looking straight ahead. After a few moments in this position, the researcher clicks the button on the Setup Window to record this posture as the reference posture for the remainder of the study.
5. The researcher triggers the beginning of the first session and leaves the participant to do their own work on the computer. The researcher remains present in case of any questions or technical issues and to let the participant in and out of the lab on breaks.
6. Feedback is disabled in the first (Baseline) session. Each session lasts 25 minutes, marked by a countdown timer (Figure 10) that appears on the computer screen. If the participant stands up the countdown pauses until they are seated again. The sensor data is recorded throughout the experiment, even while the countdown is paused.
7. After 25 minutes has passed, the first break begins. Each break lasts 5 minutes. The participant is encouraged to stand and stretch during the break, but is not required to do so. The break countdown timer (Figure 11) continues regardless of whether the participant is seated in the chair.
8. Once the first break is over and the participant is again seated in the chair, the Generic session begins. During this session the Generic feedback window (Figure 7) pops up whenever the posture threshold is exceeded. Figure 12 shows how the Generic Feedback Window appears in the context of the entire screen.

9. After the second session there is another break, identical to the first.
10. The Specific session follows the second break. During this session the Specific feedback window (Figure 8) pops up whenever the posture threshold is exceeded.
11. At the completion of the Specific session the software notifies the user that the study is finished and then closes itself.
12. The researcher then gives the participant a short paper questionnaire (Appendix A). Once the participant has completed the questionnaire the researcher thanks the participant for their time and answers any questions they might have.

Figure 10. Countdown timer for a session in the user study.

Figure 11. Countdown timer for a break in the user study.

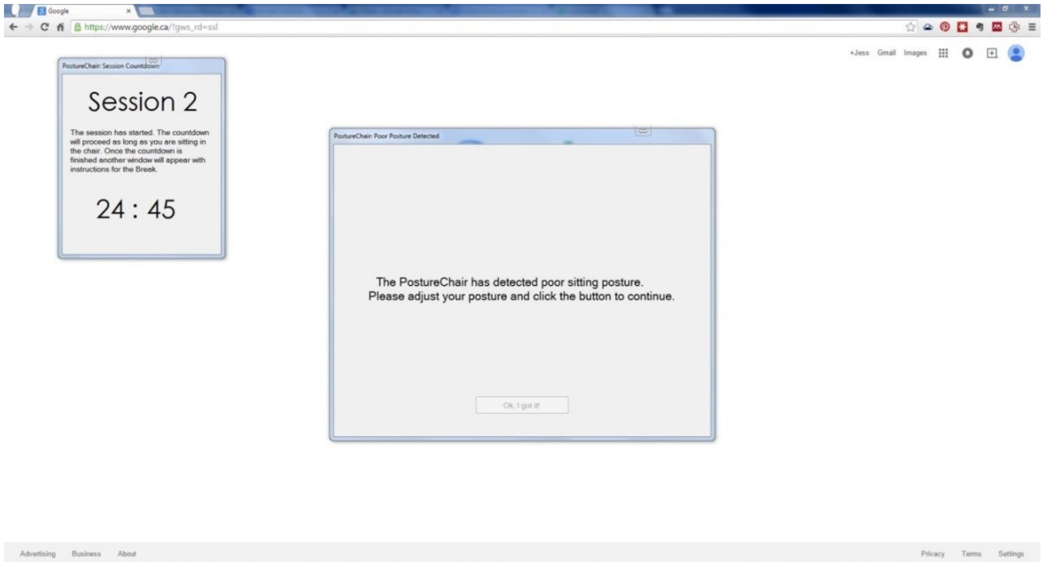

Figure 12. Participant's screen during the Generic Session once the Generic feedback window has appeared.

4.1.2 User Study Design

The user study examining the effectiveness of the PostureChair system is divided into three 25 minute sessions, with a 5 minute break between each session. This session time period was chosen for a variety of reasons. It is long enough to give users the chance to forget they are being monitored. It also provides slightly more long-term data than the 10-15 minutes studied by van der Doelen et al. [13] and Zheng et al. [62], giving a better idea of the longer-term effects of this kind of feedback on sitting posture. The session time was also influenced by the sitting time guidelines ([5, 38, 43, 45, 50]) described in Section 2.3.2. As outlined in that Section, 20-30 minutes of maximum sitting time has been found to be difficult for computer users to achieve on their own. The work of Rutten et al. [49] also suggests a pattern of 25 minutes sitting followed by a 5 minute break as a guideline for everyday work. In our user study participants were encouraged by the researcher to stand and stretch during the breaks, but they were not required to do so. In this way we attempted to break up the sitting times the participants experienced without forcing the issue as it was not the primary focus of the study. The session time was also chosen to give a total user study time of between 1.5-2 hours, depending on how often the participants stood up. We believe this was the longest user study that could be conducted without notably diminishing the number of people willing to participate.

The recruitment material for the study (Appendix B) indicated that participants could have the PostureChair system installed on their own personal laptop, so they could have access to their own documents during the study. However, the PostureChair system was designed on a Windows 7 computer and turned out to be incompatible with Windows 8 and Mac operating systems. Participants invariably brought Windows 8 or Mac laptops to the study, so it was decided to restrict the study to the networked Carleton University computers. All participants who expressed interest in the study were informed of this limitation before they came to the lab. Participants were still permitted to bring their laptops to use as a second screen or to transfer documents to the Carleton University computer mid-study, but they were instructed to use the Carleton University computer as their primary device.

4.2 Recruitment

The user study was conducted over an eight week period. Calls for participants were advertised on posters displayed on bulletin boards throughout the Carleton University campus. Calls were also posted to social media, user study participant mailing lists, and relevant online Carleton University publications (e.g. Today@ Carleton) (Appendix B). Any adult was free to participate with the exception of those with health conditions that might interfere with their ability to sit in our office chair comfortably for the entirety of the study. Participants were recruited from among the staff, faculty and graduate students at Carleton University. A few participants also came from the wider local community.

To facilitate recruitment an online form was set up to allow participants to choose an available timeslot. Each timeslot was 2 hours long and timeslots were offered over an 8 hour period for most weekdays throughout the eight week period. Participants were welcome to reschedule timeslots if necessary using the same online form.

4.3 Data Collected

Two types of data were collected during the user study: the sensor data logged by the PostureChair system and the questionnaire data collected at the end of each study. Occasional observational notes were also made over the course of each study.

4.3.1 Sensor Data

The PostureChair system logged all data to a single comma-separated text file for each study. Each log file contains the participant's ID number, the date and time the study was started, lines of sensor data and event data. The lines of sensor data include the raw data from each of the 8 sensors, the Euclidean distance calculations for each sensor, and the Euclidean distance for the overall system. The Euclidean distances are logged only after the reference posture is set. The raw sensor data for the reference posture is also logged at the time it is set. The event data lines are simply the date and time of the event and the name of the event. The events include: the start and end of the study, the start and end of each session and break, the appearance of either feedback window, the moment when a feedback window becomes dismissible, and the moment the participant clicks the button to dismiss the feedback window.

4.3.2 Questionnaire

A short paper questionnaire (Appendix A) was created to solicit participants' impression and experience of the PostureChair system, and to collect some demographic information. The questionnaire contains three types of questions. The first 15 questions are seven-level Likert scale questions, consisting of a positive or negative statement and seven choices ranging from "Strongly Disagree" to "Strongly Agree" for the participant to choose from. Questions 1-7 deal with the feedback windows and attempt to determine the participants' opinions on their disruptiveness, timeliness and helpfulness. Questions 8-10 attempt to elicit participants' opinions on their own posture, both before and after using the PostureChair system. Question 11 is standalone and addresses whether participants would be likely to use such a system in their own workplace. Questions 12-15 deal with participants' awareness of ergonomics and whether they attempt to apply that knowledge in their daily life.

Questions 16 and 17 are long-answer questions. Question 16 asks which feedback window (Generic or Specific) the participant preferred and why. Question 17 asks if there is any other information participants would have like to have seen in the feedback window. We hoped that these questions would capture any aspects of the feedback windows that we had missed in the Likert questions, and that the answers to Question 17 would give direction for future work.

Questions 18-21 are demographic questions on gender, age, height and weight. These questions are multiple choice and included the option "Prefer Not to Answer" to avoid any awkwardness participants might feel about giving specific height, weight and age data. We also hoped that giving ranges for the height and weight questions would allow participants to answer quickly and confidently, without getting caught up in attempts at unnecessary precision.

Chapter 5 Results

In total 34 participants took part in the user study: 14 male and 20 female. 9 participants (5 male and 4 female) took the study a second time in order to resolve an error in data collection, resulting in 43 data points in total. The code error, its correction, and consequences are all described in Section 5.1. Although each participant took the study alone, there were occasions when both instances of the PostureChair system were in use simultaneously in the same room. In these cases each participant was focused on their work and had little to no interaction with the other participant. Participants experienced all three of the conditions (Baseline, Generic feedback and Specific feedback) for repeated measures testing.

Participants included graduate students, university staff and a few members of the wider community. A wide range of ages was represented, as shown in Figure 13; all participants were under 70 years of age. The participants who repeated the study were a more limited subset: 8 were aged 18-29 and one was aged 30-39. There was a fairly wide distribution of body weights (shown in Figure 14) and heights among the participants. Half of the participants were between 64 inches and 69 inches tall, 9 were shorter than 64 inches and 8 were taller than 69 inches. Only one participant was taller than 74 inches and none were shorter than 58 inches. Thus the participants in this study represented a wide range of ages and body sizes, reflecting the general population of office workers.

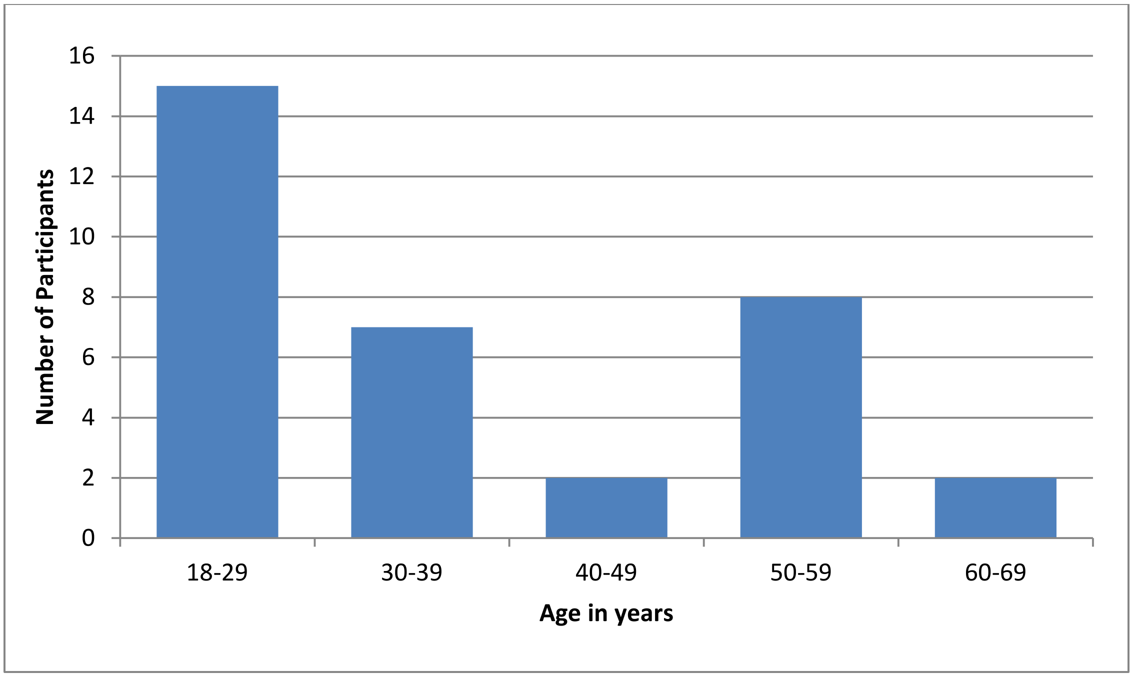

Figure 13. Distribution of participant's age in years

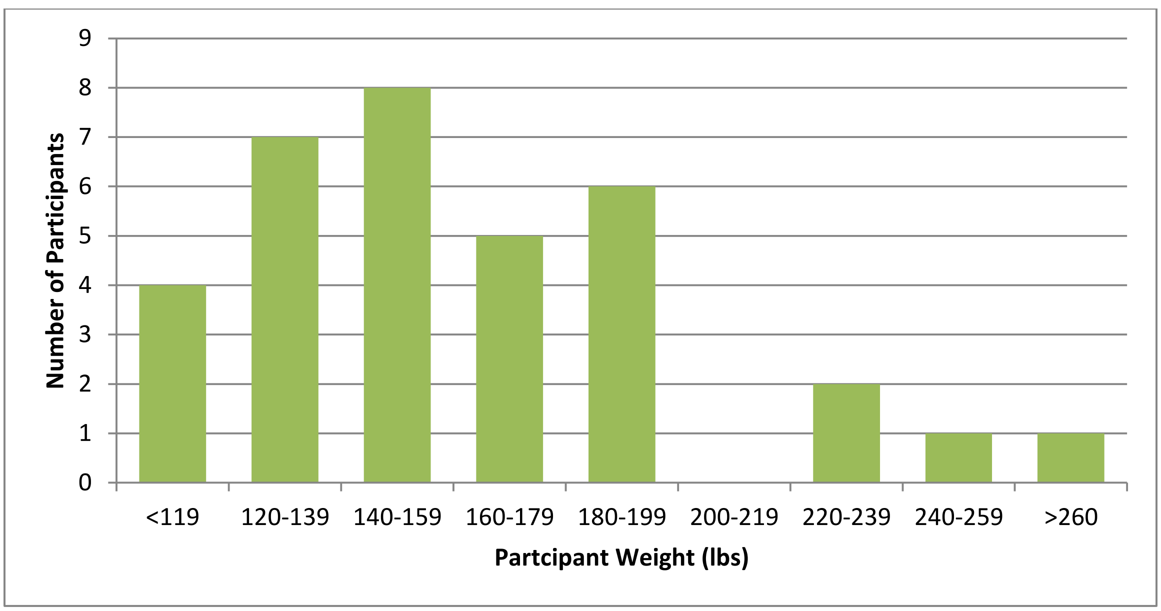

Figure 14. Distribution of participant's body weight in pounds

5.1 Error in Euclidean Distance Calculation

An error in the PostureChair software code was discovered after the first 28 participants had completed the study. This error caused the Euclidean distance calculation to incorporate only 4 of the 8 sensors in the PostureChair system. The 4 sensors that were excluded from the calculation were the 3 left-side sensors (from the user's perspective) plus the sensor in the lumbar region. Figure 15 shows the sensors from the viewer's perspective; the excluded sensors are thus on the right side and shown in red, while the sensors that were incorporated in the 4-sensor Euclidean distance calculation are shown in green. Movements on one side of the body are often reflected on the other, for instance if a user leans to the left while sitting in an augmented chair they will put more pressure on the left-side sensors and less pressure on the right-side sensors. In this way, sensors on only one side of the body can detect movement towards either direction. The sensors that were excluded in the 4-sensor Euclidean distance calculation were all along one side of the body, so only one of each pair of sensors was excluded. This led us to believe that the code error might have little effect on the posture detection and classification, which we investigate further in Section 5.1.2.

The logs from the participants who experienced the 4-sensor calculation were largely unaffected by the error since the raw data and individual Euclidean distances for each sensor were logged for each sensor reading. As a result the logs were updated with the correct Euclidean distance calculation for the entire 8-sensor system. However, the code error did affect how often the feedback windows appeared to each of the first 28 participants. Since the code error was in place during pilot testing the posture threshold that was selected was based on this 4-sensor Euclidean distance calculation.

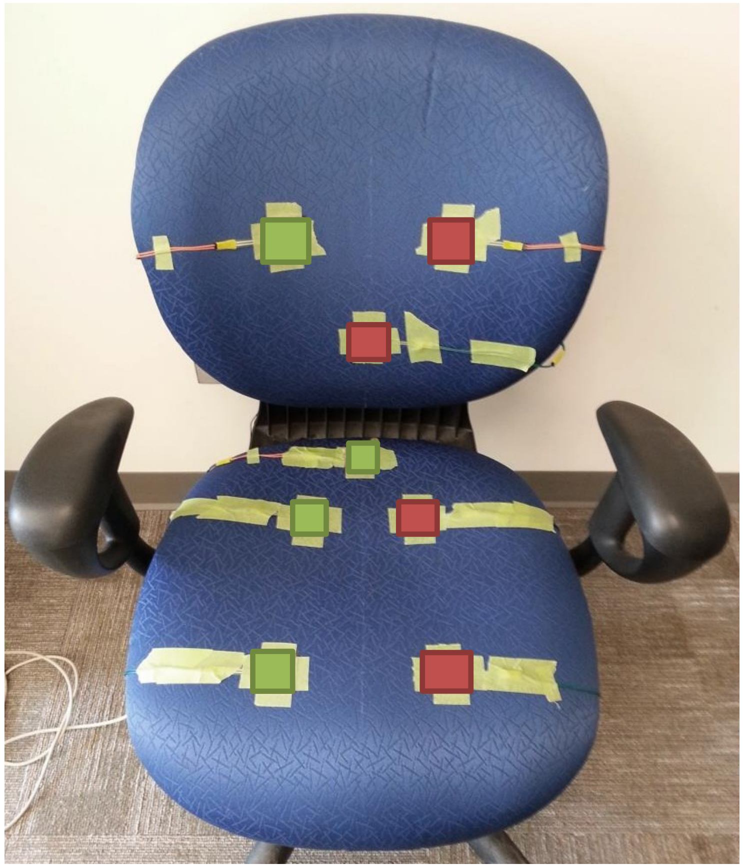

Figure 15. Augmented chair showing which FSRs were included in the 4-sensor Euclidean distance calculation (shown in green), and which were excluded (shown in red)

The error was corrected in the software as soon as it was detected and from then on all 8 sensors were incorporated into the overall Euclidean distance calculation. The pilot study data was re-calculated using the 8-sensor Euclidean distance calculation and a new posture threshold was selected through the same process as the initial, 4-sensor posture threshold. The time thresholds were not changed. The PostureChair software was then updated with the new posture threshold. Following this correction, 9 participants were invited to repeat the study; for most of these participants about 4-6 weeks passed between their first and second user study. In addition, 6 new participants were found to take the study with the error corrected. The logs for the participants who completed the study before the error was found were updated with a re-calculated 8-sensor Euclidean distance. All posture data discussed and analyzed in this Chapter is based on this 8-sensor calculation. The only difference in the user study before and after the error was corrected is the context that influenced how many feedback pop-ups the participants actually saw.

5.1.1 User Study Groups

In total, 43 data points were collected: 28 participants took part in the study with the code error in place and 15 participants took part after the error was corrected. The participants can therefore be divided into two groups: the 4-sensor Euclidean distance group and the 8-sensor Euclidean distance group. This can be described in set notation:

$$\{\text{Total User Study}\} = \{\{\text{4-Sensor}\}, \{\text{8-Sensor}\}\}$$

$$|\text{Total User Study}| = 43$$

These subsets of participants can be further broken down by the number of times each participant took part in the study. 19 of the {4-Sensor} participants and 6 of the {8-Sensor} participants took part in the study once ("Single Run"), while 9 participants took part in the study twice, experiencing both the 4- and 8-sensor calculations ("Double Run"). In set notation:

$$\{\text{4-Sensor}\} = \{\{\text{Single-Run \#1}\}, \{\text{Double-Run \#1}\}\}$$

$$\text{where } |\text{Single-Run \#1}| = 19, |\text{Double-Run \#1}| = 9$$

and

$$\{\text{8-Sensor}\} = \{\{\text{Single-Run \#2}\}, \{\text{Double-Run \#2}\}\}$$

$$\text{where } |\text{Single-Run \#2}| = 6, |\text{Double-Run \#2}| = 9$$

In total there were 25 single-run participants and 9 double-run participants, which is 34 participants in total. Of the set of all participants, 4 were found to have invalid data due to malfunctioning sensors. The malfunctioning sensors occurred with only single-run participants who had experienced the 4-sensor calculation:

$$\{\text{Single-Run \#1}\} = \{\text{Valid Single-Run \#1}\} + \{\text{Invalid}\}$$

$$\text{where } |\text{Valid Single-Run \#1}| = 15, |\text{Invalid}| = 4$$

All 4 of these invalid data points were excluded from all analyses. In the rest of this Chapter whenever the {Total User Study}, {4-Sensor} or {Single-Run #1} sets are referred to they can be understood as excluding the {Invalid} subset.

5.1.2 Effect on Results

To determine if the code error significantly affected the experience of the first 24 valid 4-Sensor participants, the total number of feedback pop-ups seen in the course of each study was compared between the 4-Sensor and 8-Sensor groups. The data for the 4-Sensor group and the 8-Sensor group were both found to be significantly non-normal by the K-S test, as shown in Table 1. There was no significant difference found between the number of pop-ups seen by the 4-Sensor group and the 8-Sensor group when compared using a non-parametric Mann-Whitney test, $U = 157.0$, $z = -0.67$, $p = .521$.

We also compared the total number of feedback pop-ups seen by the 9 Double-Run participants. The total number of feedback pop-ups seen by the Double-Run participants in the 4-Sensor group was found to have a normal distribution using the K-S test. However, the total number of feedback pop-ups seen by the Double-Run participants in the 8-Sensor group was found to be significantly non-normal, as shown in Table 1. There was no significant difference found between the number of pop-ups seen by the Double-Run participants when they experienced the 4-sensor calculation and when they experienced the 8-sensor calculation, when compared using a non-parametric Mann-Whitney test, $U = 39.0$, $z = -0.13$, $p = .931$.

Together, these two results suggest that the code error likely did not create a significantly different experience for the participants in the 4-Sensor group than they would have experienced if there had been no error. Consequently we treat the posture data from participants in both the 4-Sensor and 8-Sensor groups as equivalent in the following Sections.

Table 1. Descriptive statistics and K-S test results for total number of feedback popups that appeared over an entire user study, by user study subset. Significant results shown in bold.

	Descriptive				K-S	
	N	*Mean*	*Std Dev*	*Mdn*	*D*(N)	*p*
Data for both Single-Run and Double-Run participants						
4-Sensor Euclidean	24	16.8	15.1	11.5	0.229**	**.002**
8-Sensor Euclidean	15	20.3	15.7	16.0	0.230*	**.032**
Data for only Double-Run participants						
4-Sensor Euclidean	9	21.1	17.0	15.0	0.196	.200[a]
8-Sensor Euclidean	9	20.1	17.7	9.0	0.291*	**.027**

*$p < .05$; **$p < .01$, [a]This is a lower bound of true significance

5.2 Improvement in Posture

This study addressed two main research questions:

1. Can contextual, digital feedback improve the sitting posture of computer users?
2. Does this feedback need to include specific instructions on how to improve posture in order to be effective?

To quantify improvement in posture, we define bad posture as any 8-sensor Euclidean distance calculation above the posture threshold. Each 25 minute session can then be scored by the fraction of time a participant spent in bad posture in that session.

5.2.1 Posture Variation between Participants

One of the challenges of designing a system that detects and attempts to modify sitting posture is the extreme variation in behaviour between users. The data from the participants in this study reflects this variation: a few participants spent almost no time in a bad sitting posture, and some spent one or more of the sessions entirely in bad posture. Figure 16, Figure 17, and Figure 18 show the 8-sensor Euclidean distance results over all three user study sessions for 3 different participants. Figure 16 is the graph of a participant who had mostly good posture and saw only 2 feedback pop-ups. Note that the plot remains under the posture threshold for almost the entire period under study, except during the breaks when the participant stood up. The smaller spikes in the plot during the study sessions are due to the participant shifting in the chair and momentarily removing pressure from one or more sensors.

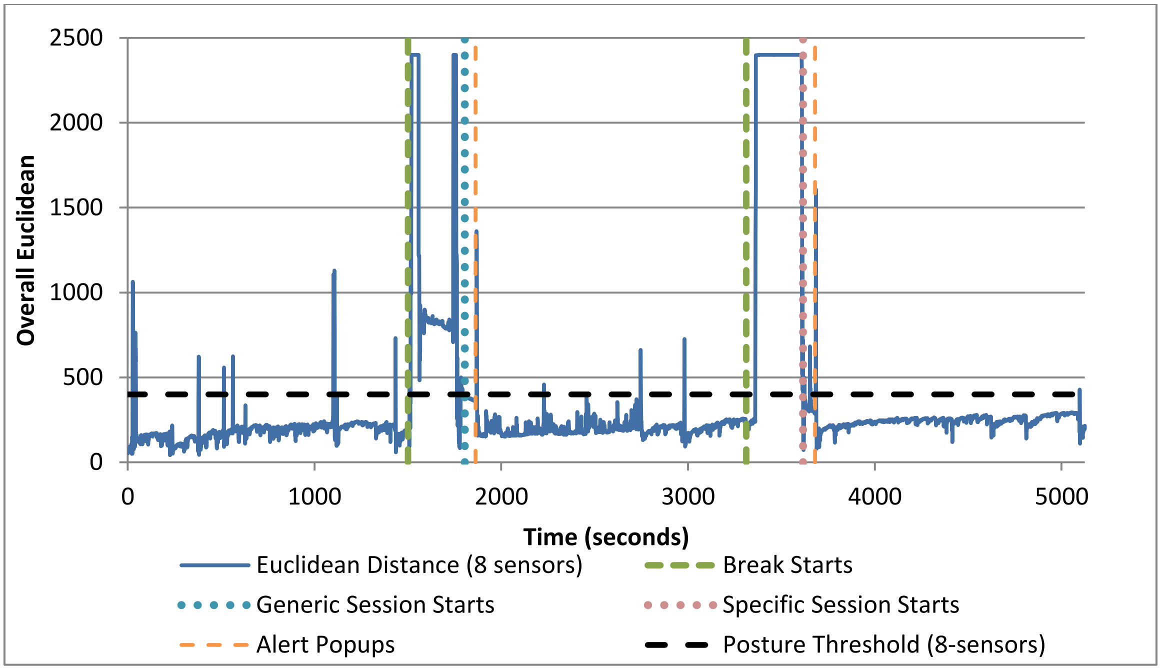

Figure 16. 8-sensor Euclidean distance data for Participant 22 over the entire user study. The strong dashed line in black represents the posture threshold, and the dashed green lines represent the start of each break. Dotted lines represent the start of each session, and the dashed orange lines represent the appearance of the feedback pop-ups.

Figure 17 shows the polar opposite: a participant who spent all three sessions in bad posture. Note that this plot never dips below the posture threshold and that the participant remained seated during both breaks. This plot is not a case of sensor malfunction and this participant experienced the 8-sensor Euclidean distance calculation. The reference posture was examined for any error at the time of the recording and none was found. The participant simply adopted their usual sitting posture after the reference posture was set and never returned to the reference posture. This indicates one of the limitations of the PostureChair system: participants can simply choose to ignore the feedback entirely. This is a limitation of any persuasive system. Perhaps if the user could customize the posture threshold, lowering the distance between the threshold and the participant's usual posture, the system could affect greater change in the participant's posture. However, it is also possible that this participant was simply not responsive to digital feedback and that a different method of persuasion would be required to effect change in this case.

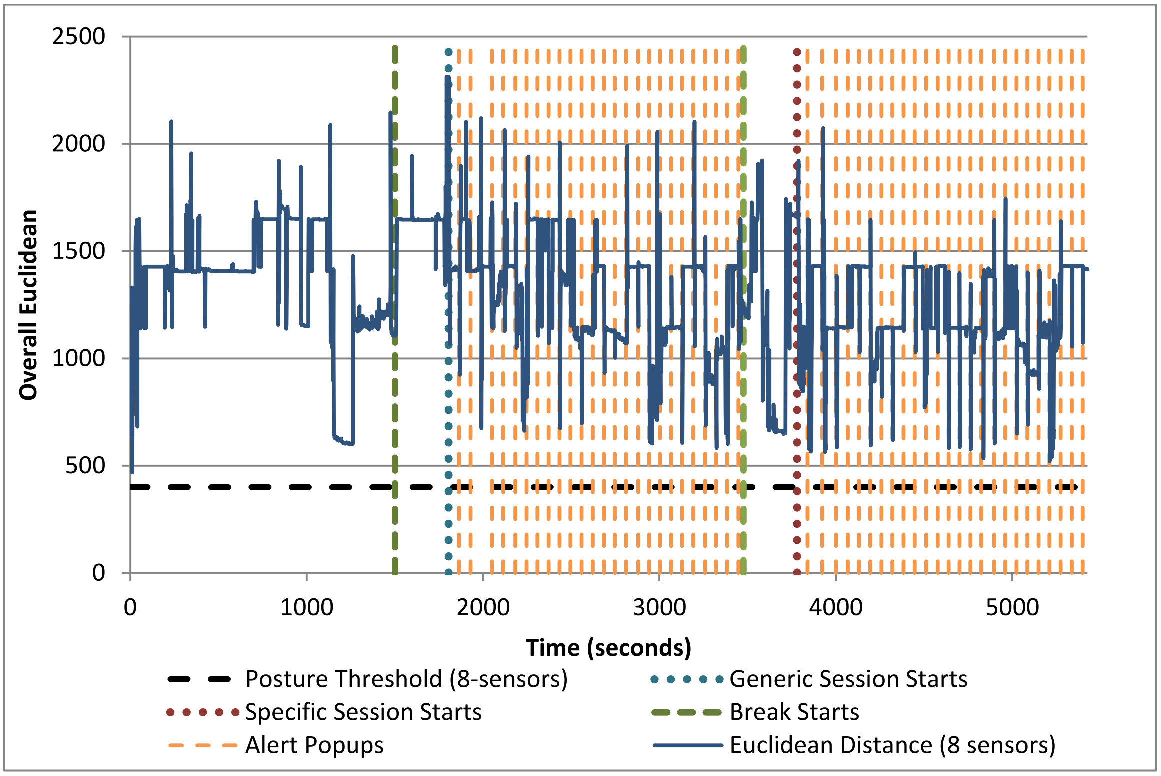

Figure 17. 8-sensor Euclidean distance data for Participant 35 over the entire user study. The strong dashed line in black represents the posture threshold, and the dashed green lines represent the start of each break. Dotted lines represent the start of each session, and the dashed orange lines represent the appearance of the feedback pop-ups.

Figure 18 shows a more typical posture plot. This participant spent considerable amounts of time in both good and bad posture, and the amount of time they spent in good posture increased in the feedback-enabled sessions. There is also a drop in the plot after each feedback pop-up appears, suggesting that the participant responded to each pop-up and moved closer to the reference posture. The participant stood for at least part of each break. This plot represents the ideal case for the PostureChair system: a user that does not already have near-perfect posture, and who responds well to the digital feedback provided by the system. In the non-ideal cases, it is our opinion that a more customizable system could show better results, which is a possible direction for future work.

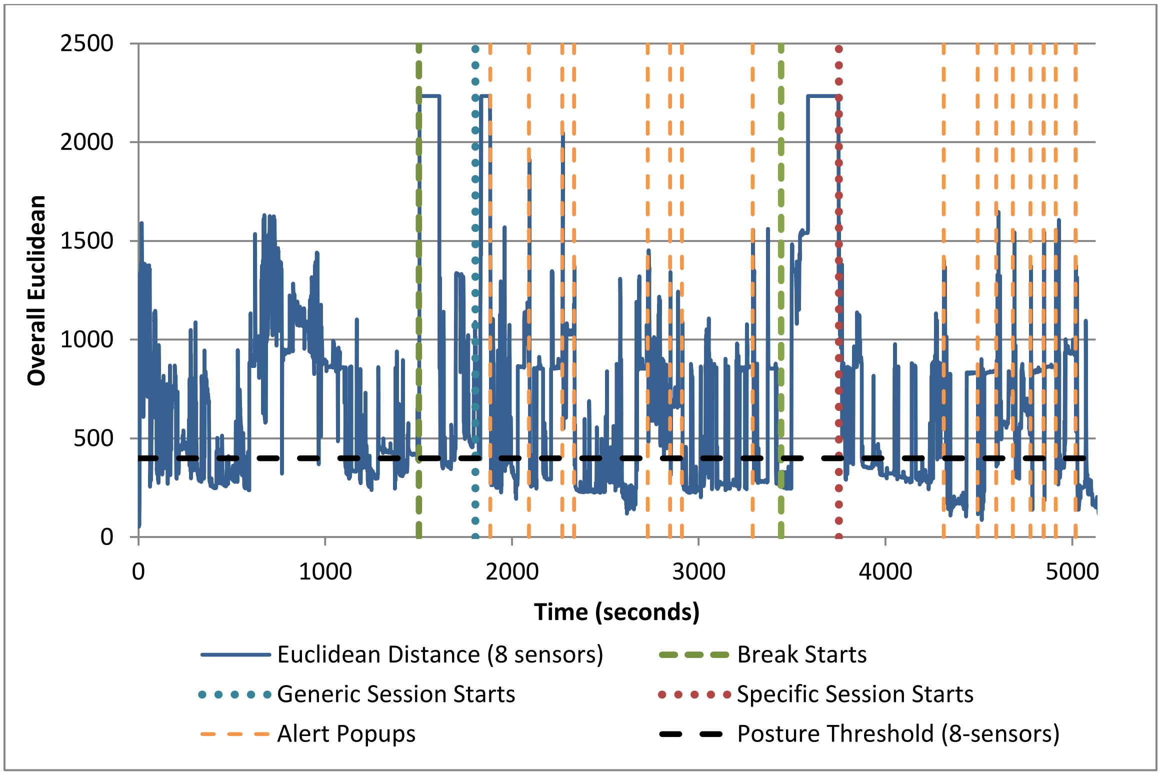

Figure 18. 8-sensor Euclidean distance data for Participant 40 over the entire user study. The strong dashed line in black represents the posture threshold, and the dashed green lines represent the start of each break. Dotted lines represent the start of each session, and the dashed orange lines represent the appearance of the feedback pop-ups

5.2.2 Posture and Demographics

Throughout the user study we observed that the fixed nature of the office chair and desk was a distinct limitation of the system, since for some participants one or the other simply did not match their body size or shape. For instance, several participants found the desk height too high; they could set the chair at the right height for their legs or for ergonomic keyboard use, but not both. We included a footrest to accommodate this issue, but a few participants preferred to rest their feet on the chair wheels instead. Other participants found that the lumbar curve did not match the shape of their spine. A few participants found the chair too small and one or two expressed a preference for a higher desk height. A more adjustable chair would have helped this issue, but the ideal solution would be a variety of augmented chairs of different sizes. Whether these discrepancies between the chair size, desk height, and the participants' body shapes affected the reliability of the chair is unknown. However, this conflict does match the situation in most workplaces, where chairs are generally one-size-fits-all. In this way the PostureChair system and our study accurately reflect the realities of office work environments.

This observation of the discrepancy between chair size and participant size led to an examination of potential correlations between posture and participant demographic data such as gender, age, height and weight. In the following statistical tests only the fraction of time spent in bad posture for the Baseline session was considered as it was the only part of the study in which the participants' posture was unaffected by feedback. Also, since there were only 1 or 2 participants in the higher weight categories, these were combined into a single category of greater than 200 lbs. The subset of participants under consideration here is the set {{Total User Study} – {Double Run #2}}, 30 participants in total, since the demographic data is repeated for Double Run #1 and Double Run #2.

Table 2 shows the descriptive statistics and the results of the K-S test for normality for the variables under consideration. The age, height and weight variables were recorded in categories, so no mean was calculated. Gender was considered as a binary variable and so no test for normality was necessary. As shown in this table all of the variables are significantly different from a normal distribution, so the non-parametric Spearman's rho was used as the correlation statistic.

Table 2. Descriptive statistics and K-S test results for Baseline posture and participant age, height and weight. Medians are described using the category label. Significant results shown in bold.

	Descriptive			K-S Test	
	Mean	*Std Dev*	*Mdn*	*D*(30)	*p*
Baseline Posture	0.56	0.37	0.63	.184*	**.011**
Age	n/a	n/a	30-39 years	.281**	**.000**
Height	n/a	n/a	64"-69"	.281**	**.000**
Weight	n/a	n/a	140-159 lbs	.171*	**.025**

*$p < .05$; **$p < .01$

The fraction of time spent in bad posture in the Baseline session was not significantly correlated with participant gender, age, weight or height in this small sample. Whether a correlation between these variables would be found with a larger sample size is still an open question.

5.2.3 Posture Results by User Study Group

There are multiple ways to group the participants in this user study, as laid out in Section 5.1.1. As some of the differences are based on how many times a participant took part in the study and how the Euclidean distance was calculated, this Section will include comparisons of bad posture times within each subset of participants. Besides those sets laid out in Section 5.1.1, two other subsets are defined for the comparisons in this section: {{Total User Study} – {Double Run #1}} excludes the first user study experience of each Double-Run participant, and {{Total User Study} – {Double Run #2}} excludes the second user study experience of each Double-Run participant. These subsets were included in the testing to determine if the data from the Double-Run participants was affected by those participants taking part in the study twice. Recall that the {4-Sensor} and {8-Sensor} subsets refer to how many sensors were involved in the Euclidean distance calculation during the user study. All data in the 4-Sensor group was subsequently re-calculated using all 8 sensors, and it is this re-calculated data that is discussed here. All tests were carried out with the {Invalid} set excluded.

Table 3 shows the number of participants in each subset, as well as the mean fraction of time the participants in each subset spent in bad posture in each session. Participants generally spent less time in bad posture in the Generic session than in the Baseline session, and even less time in bad posture in the Specific session, which can be seen in Figure 19.

Table 3. Descriptive statistics for the fraction of time spent in bad posture in each session, for each set of participants

		Baseline		Generic		Specific	
Subset of Participants	N	*Mean*	*Std Dev*	*Mean*	*Std Dev*	*Mean*	*Std Dev*
{Total User Study}	39	0.62	0.35	0.47	0.29	0.42	0.33
{Total User Study} – {Double Run #2}	30	0.56	0.37	0.44	0.30	0.41	0.33
{Total User Study} – {Double Run #1}	30	0.61	0.35	0.47	0.30	0.44	0.33
{4-Sensor}	24	0.56	0.38	0.42	0.30	0.39	0.33
{8-Sensor}	15	0.72	0.29	0.54	0.27	0.47	0.34

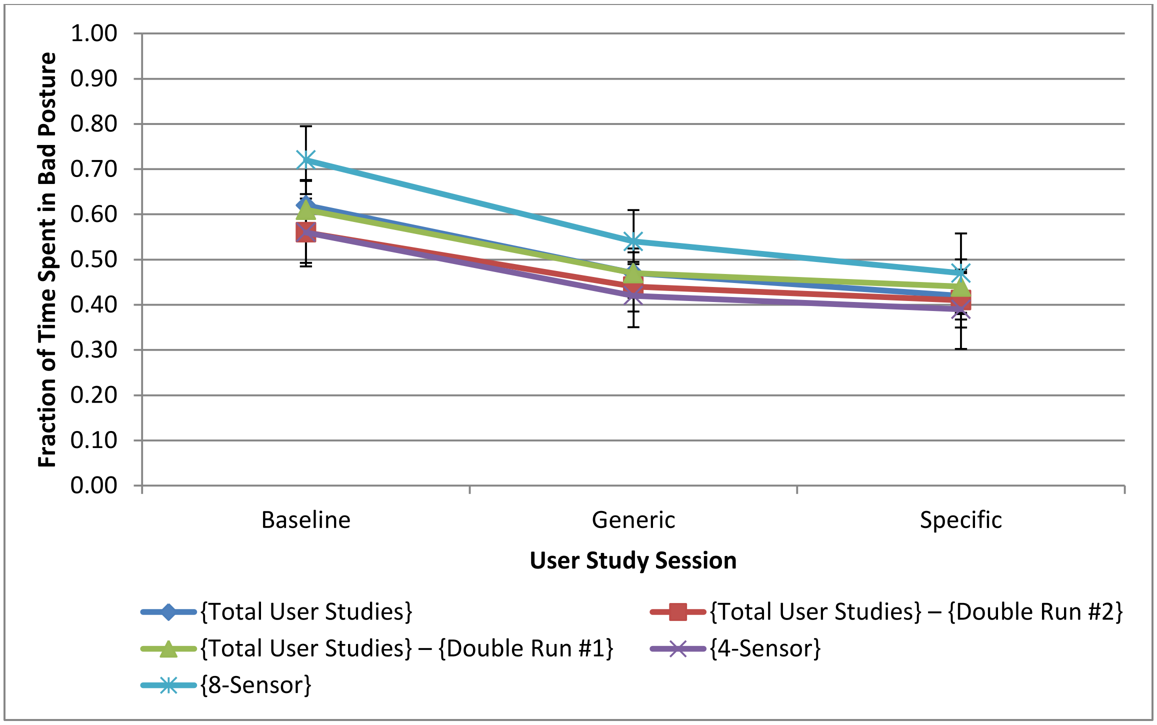

Figure 19. Fraction of time spent in bad posture in each user study session, for each set of participants. Error bars show standard error.

The fractions of time participants spent in bad posture in each of the three sessions (Baseline, Generic Feedback and Specific Feedback) were compared using a repeated-measures ANOVA for each of the subsets. Table 4 shows that for 4 of the 5 subsets Mauchly's test indicated that the assumption of sphericity had been violated, so degrees of freedom for those subsets were corrected using the Greenhouse-Geisser correction of sphericity. The 8-Sensor subset did not violate the assumption of sphericity and therefore no correction was made to the degrees of sphericity. Table 4 also shows that bad posture was significantly affected by the feedback type (None, Generic or Specific) in every subset of participants except {4-Sensor}, which was approaching significance. Recall that some data points belong to more than one subset, so members of the 4-Sensor subset belong to every other set described in Table 4 except the 8-Sensor subset. Also, the 4-Sensor participants did not experience a significantly different number of feedback pop-ups than the other participants, as shown in Section 5.1.2. This suggests that the lack of significant effect of feedback type in the 4-Sensor set is more likely due to variation between participants rather than the code error.

Table 4. Repeated measures Analysis of Variance: statistics for Mauchly's test and effects of feedback type. Significant results shown in bold.

	Mauchly's Test			ANOVA		
Subset of Participants	$\chi^2(2)$	p	ε	F	df	p
{Total User Study}	17.32**	**.000**	.73	8.478**	(1.47, 55.32)	**.002**
{Total User Study} – {Double Run #2}	9.92**	**.007**	.77	3.839*	(1.24, 44.68)	**.039**
{Total User Study} – {Double Run #1}	8.62*	**.013**	.79	6.629**	(1.58, 45.85)	**.005**
{4-Sensor}	15.72**	**.000**	.66	3.582	(1.32, 30.45)	.057
{8-Sensor}	2.65	.265	-	5.155*	(2, 28)	**.012**

*p < .05; **p < .01

The repeated-measures ANOVA also included contrasts between each of the sessions; the results are shown in Table 5. Together with the descriptive statistics in Table 3 these results show that the participants spent significantly less time in bad posture in both of the feedback enabled sessions (Generic and Specific) compared to the no feedback session (Baseline). This significant difference occurred in every subset under consideration except the 4-Sensor subset, although even the difference in the 4-Sensor subset was approaching significance. These results are consistent with the overall effects of feedback type described above. The difference in bad posture times between the Generic and Specific sessions is less clear. Participants in every subset spent less time in bad posture in the Specific session than in the Generic session, however this difference is significant for only 3 of the 5 sets: {Total User Study}, {{Total User Study} – {Double Run #1}}, and {8-Sensor}. All three of these sets showed a more significant main effect of feedback type than the other two sets.

Table 5. Repeated measures Analysis of Variance: statistics for within-subjects contrasts between feedback types. Significant results shown in bold.

	Overall	Baseline vs. Generic		Baseline vs. Specific		Generic vs. Specific	
Subset of Participants	*df*	*F*	*p*	*F*	*p*	*F*	*p*
{Total User Study}	(1, 38)	9.398**	**.004**	10.035**	**.003**	7.639**	**.009**
{Total User Study} – {Double Run #2}	(1, 29)	4.414*	**.044**	4.657*	**.039**	3.310	.079
{Total User Study} – {Double Run #1}	(1, 29)	8.135**	**.008**	7.973**	**.008**	5.221*	**.030**
{4-Sensor}	(1, 23)	4.245	.051	3.913	.060	2.971	.051
{8-Sensor}	(1, 14)	5.261*	**.038**	6.953*	**.020**	5.058*	**.041**

*$p < .05$; **$p < .01$

This is a difficult set of results to parse, made more complicated by the number of different subsets to consider. However, since the 4-Sensor set is a subset of the Total User Study set, and the Total User Study set showed a significant effect of feedback type, it is likely that the lack of significant effect in the 4-Sensor set is due to variation between participants, which has a greater impact on the smaller set. Recall that there is no overlap between participants in the 4-Sensor subset and the only smaller set, the 8-Sensor subset. Overall, both types of feedback had a significant effect on participants' sitting posture when compared to no feedback. The effect of the Specific feedback on posture when compared to that of the Generic feedback is less clear. The Specific feedback is significantly better at improving participants' posture than the Generic posture to $p < .01$ for the set that encompasses all data points (Total User Study). However, this improvement is significant to a lesser degree ($p < .05$) or not significant at all in the subsets. We posit that the information found on the Specific feedback pop-up is only useful to some participants, and not others, while the feedback pop-ups in general are effective for most participants. This would explain the lesser effect and wider variation between subsets.

Returning to the research questions, these results show that contextual, digital feedback can significantly improve computer users' sitting posture. The simple contextual feedback found in the Generic pop-up is sufficient to produce an improvement in posture, and the additional information given in the Specific feedback window is not strictly necessary. Any additional improvement in posture created by the Specific feedback above that created by the Generic feedback appears only in the largest set of participants and in certain subsets. However, the additional information may affect the user's experience with the PostureChair system, even if it does not affect their actual posture. Results related to the participants' experience with the feedback windows are covered in Section 5.3. It is possible that the prior knowledge of the Double Run participants could have affected the results of their second experience of the study, so all results in Section 5.3 and 5.4 are examined using data from the {{Total User Study} – {Double Run #2}} set.

5.2.4 Implications for Chair Design

As a result of the error in the PostureChair software described in Section 5.1, 28 participants completed the study with a system in which the feedback behaved as if only 4 sensors were present on the chair. The data from all 8 sensors were logged to the log file, and the only feature that used all 8 sensors was the sensor indicators on the Specific feedback pop-up. These still functioned as intended and reflect the status of all 8 sensors. Despite the code error, we have shown that there was a lack of a significant difference in the number of pop-ups participants saw with the 4-sensor and 8-sensor calculation. While this does not definitively prove that there was no significant different in participant experience in these two conditions, it does indicate that this is likely the case. There may be some difference in the effect of the feedback pop-ups on the participants' posture, but the evidence of this is less clear. As shown in Section 5.2.3, there was no significant effect of feedback type on posture for the 4-Sensor set. However, all the other sets, some of which overlap with the 4-Sensor set, did show a significant effect of feedback type on posture. These results lead us to conclude that the difference in the 4-Sensor set is more likely due to variation among participants rather than a property of the 4-Sensor posture detection. This conclusion requires further investigation before it can be stated with proper assurance.

Most recent papers have settled on 6-8 sensors [27, 62] as sufficient to detect and classify sitting posture as one of 10 possible postures. Any works with fewer sensors typically only attempt to recognize a directional lean rather than actual posture [10]. However, the results of the 4-Sensor setup in the PostureChair system suggest that as few as 4 sensors, placed on one side of the user's body, could be sufficient when attempting a binary “good” vs “bad” classification of posture. This binary classification also appears sufficient to create effective Generic feedback to the user. Even though a single FSR costs around $7-$10 retail, which is not expensive, using fewer sensors would lower the cost of producing a commercial augmented office chair. However, all 8 sensors were used in the Specific feedback graphic, and it is likely that any feedback that attempts to give more than a simple notification of bad posture would require data from more than 4 sensors.

As discussed previously, the 4 sensors included in the 4-sensor calculation were all along one side of the body. Since movements on one side of the body are often reflected on the other, sensors on only one side of the body can detect movement towards either direction. The only central sensor that was missing from the 4-Sensor setup was in the lumbar region. Anecdotally, we observed that the lumbar region of the office chair often did not match the lumbar curve on the participants’ back. This “mismatch” appeared to occur more frequently than any other issue with the chair size or shape. The 4-Sensor results suggest that this sensor may not be crucial for a binary posture classification, although it may be necessary for more detailed classification. It is also possible that a lumbar sensor may be more helpful with a more adjustable chair that would place the sensor in the actual lumbar region of the user’s back. The lumbar data from the current office chair may be lost in the noise generated by participants for whom the lumbar support was in a less-than-ideal position. Thus, although we have shown that the 4-Sensor setup appears to be sufficient for binary posture classification and Generic feedback, we believe more investigation is necessary to confirm whether the lumbar sensor should be included or not.

This comparison of a 4-Sensor and an 8-Sensor setup was not a part of the original research question for this project. Instead, it was the result of an accident that allowed us to investigate a new aspect of our system. A more rigorous comparison is necessary to state with confidence that 4 sensors are sufficient for binary posture classification in an augmented chair system. However, these results do suggest that for a binary classification far fewer sensors may be needed than are currently predominant in the literature. They also support Mutlu et al.'s [37] assertion that the location of the sensors is more important than the absolute force measured at those sensors for posture classification.

5.3 Experience of Digital Feedback

Although the main focus of this research project is to determine if digital feedback pop-ups are effective at changing participants' posture, how participants experienced the feedback is also an important factor in its success. If feedback is disruptive or if users find it irritating, then they will be less likely to use the PostureChair system even if it is effective at changing their posture. The feedback in our system was designed to be moderately disruptive to ensure that participants would not miss or ignore any of the feedback pop-ups, which would have confounded subsequent analysis. Examining the disruptiveness of the feedback and how participants reacted to it will give some indication if a future version of the system should be re-designed with more appealing feedback. Any redesign would have to be examined to ensure that it remained effective at improving participants' posture, since that is the primary goal of the system.

As stated above, it is possible that the prior knowledge of the Double Run participants could have affected the results of their second experience of the study, so all results in Section 5.3 and 5.4 are examined using data from the {{Total User Study} – {Double Run #2}} set.

5.3.1 Disruptiveness of Feedback Windows

The disruptiveness of the feedback pop-ups was not the main focus of study, so no workload assessments or other quantitative measures of disruptiveness were carried out. Instead, the amount of time each feedback window was open was recorded since the longer the feedback pop-ups are open, the more likely they are to be disruptive. Participants were also asked if they found the feedback pop-ups disruptive and if they thought the pop-ups appeared too often on the post-study questionnaire.

On two occasions a participant stood up out of the chair during one of the feedback sessions and the feedback pop-up was left open for an extended period of time. These instances were removed from the analysis since the amount of time the feedback pop-up was open was not related to the participant's reaction to the feedback but rather their absence from the computer. Excluding these two instances, the total amount of time feedback pop-ups were open in the Generic and Specific feedback sessions was calculated for each participant. A K-S test on the difference between the two totals for each participant was significant, $D(30) = 0.208$, $p = .002$, so the differences are significantly different from a normal distribution. As a result, the totals were compared using the non-parametric Wilcoxon signed-rank test. For 12 participants the feedback pop-ups were open longer in the Generic session, and for 16 participants they were open longer in the Specific session. 2 participants saw both feedback pop-ups for an equal amount of time over both sessions. There was no significant difference between the amount of time participants saw the Generic pop-up compared to the Specific pop-up, T = 190.5, $p = .784$. This suggests that overall participants did not take more time to look at the extra information on the Specific pop-up than they did to register the appearance of the Generic pop-up. The Specific feedback is thus likely not any more disruptive than the Generic feedback.

This lack of significant difference between the total times the pop-ups were open may have been affected by the very short times the pop-ups were open towards the end of each session. The first Generic pop-up to appear in the session was only open on average 8.9 seconds and at most 21 seconds, while the first Specific pop-up was open on average 32.3 seconds and at most 363 seconds. The Specific pop-up that was open for 363 seconds seems to be an outlier, with the next longest pop-up open for 82 seconds. Removing this outlier reduces the average time the first Specific pop-up was open to 20.5 seconds, which is still more than twice the average time the first Generic pop-up was open. Participants closed both types of feedback pop-ups more quickly the more pop-ups they had already seen, as demonstrated in Figure 20. This Figure also shows the number of participants that saw each number of pop-ups, and the different trends for the Generic and Specific pop-ups. The Generic pop-ups follow a roughly linear trend, while the Specific pop-ups start with a high initial average opening time that decreases dramatically at first, then more linearly after. As the graph shows, the average opening times for pop-ups that came later in the session are based on fewer participants than those that came earlier in the session. In addition, those who dismissed the feedback windows quickly may not have been adjusting their posture as a result, which would have led them to see higher numbers of pop-ups than those who adjusted back to the reference posture.

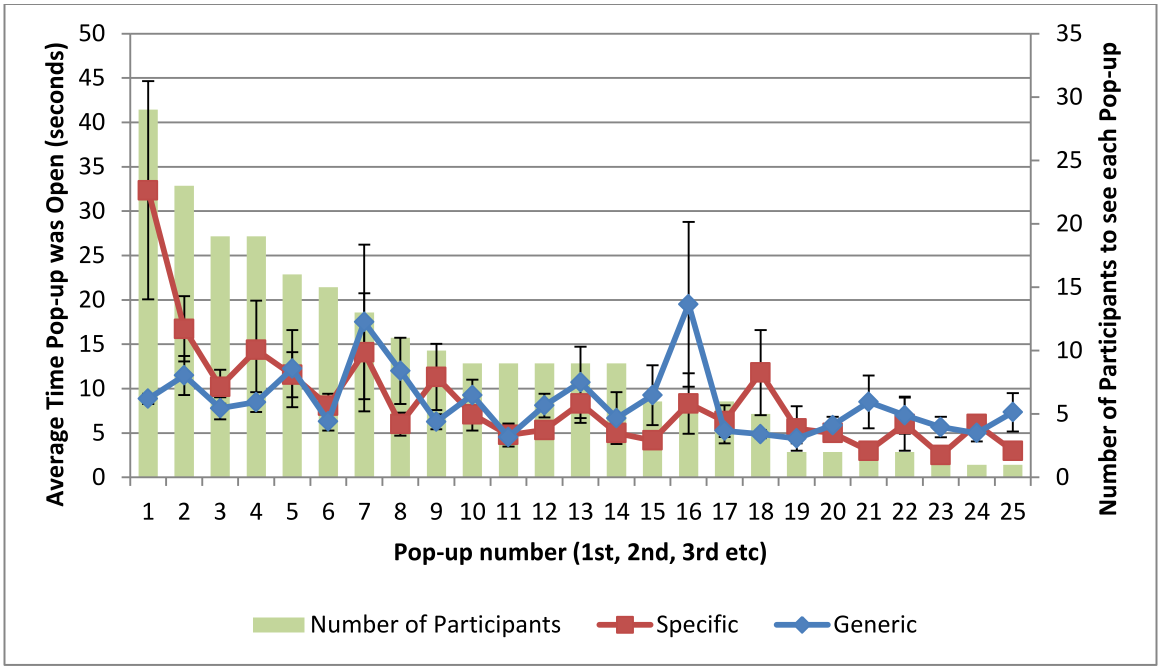

Figure 20. Mean time each feedback pop-up was open in the Specific and Generic sessions. Error bars show standard error

It is likely that the Specific pop-ups were open longer because of the extra information they presented. Several participants were observed experimenting with the Specific pop-ups by moving around in the chair to see what would cause the circles to turn red or grey. This may also explain the high average opening time for the first Specific pop-up as participants likely experimented only with the first few pop-ups they saw. Participants were shown a photo of the Specific feedback pop-up at the introduction to the study but did not experience the circle colours updating dynamically until the first pop-up appeared on their computer screen.

As part of the post-study questionnaire participants were asked how much they agreed with the statement "The feedback pop-ups were disruptive to my workflow", on a 7-point scale Likert from "Strongly Disagree" to "Strongly Agree". Descriptive statistics and the results of the K-S normality test for the responses to all 15 Likert questions on the questionnaire can be found in Appendix D. All of the responses were found to be significantly different from a normal distribution, and all correlations involving the responses are reported using the non-parametric Spearman's rho statistic. The median score on the disruptiveness question was 3, or "Slightly Disagree", so less than half of the participants found the pop-ups even slightly disruptive. Participants were also asked whether they agreed with the statement "The feedback pop-ups appeared too often", on the same scale. The median score for this question was 2.5 (halfway between "Slightly Disagree" and "Disagree), so a majority of participants thought the frequency of the pop-ups was acceptable. Participants saw a median of 13 feedback pop-ups over the entire user study.

The responses to these two Likert questions and the total number of pop-ups seen were all found to be significantly correlated. Participants who found the feedback pop-ups to be disruptive also thought the pop-ups appeared too often, $r_s = .587$, $p = .001$. This could explain why later pop-ups were open for less time than the first few: participants likely found these later pop-ups more disruptive and dismissed them more quickly in order to return to their main task. The nature of digital feedback could also contribute to the disruptiveness but this project focuses only on a single feedback type and is unable to answer that question. More importantly, participants who saw more feedback pop-ups were more likely to find them disruptive, $r_s = .642$, $p = .000$, and to think they appeared too often, $r_s = .654$, $p = .000$. This supports our assertion that a real-world version of the system should be made customizable. If users could alter the frequency of the pop-up appearances and the posture threshold, they could prevent undue disruption of their work. This customization could cause the system to be less effective, however it would also likely make the system more appealing and more likely to be adopted by users whose posture requires considerable adjustment.

5.3.2 Participant Opinion on Feedback Accuracy

There were three questions in the questionnaire (Appendix A) related to the accuracy of the feedback pop-ups. Two of these statements refer to the system overall and one refers to Specific feedback alone. Participants were asked how much they agreed with the following statements:

- “The chair accurately identified when I had good posture”
- “The feedback popups appeared in a timely manner”
- “The specific instructions on how to change my posture matched my actual posture at the time”

As shown in Figure 21, most participants agreed with each of these statements. Overall, participants indicated that the chair accurately detected their good posture and that the feedback appeared quickly after detecting bad posture (i.e. “in a timely manner”). They were less sure about the instructions on the Specific feedback pop-up, with 6 participants feeling to some extent that the instructions did not match their actual posture. This indicates that there is some room for improvement with the Specific feedback pop-up, and that perhaps more work should be done with the single sensor posture threshold to ensure that the Specific feedback reflects the user’s actual posture.

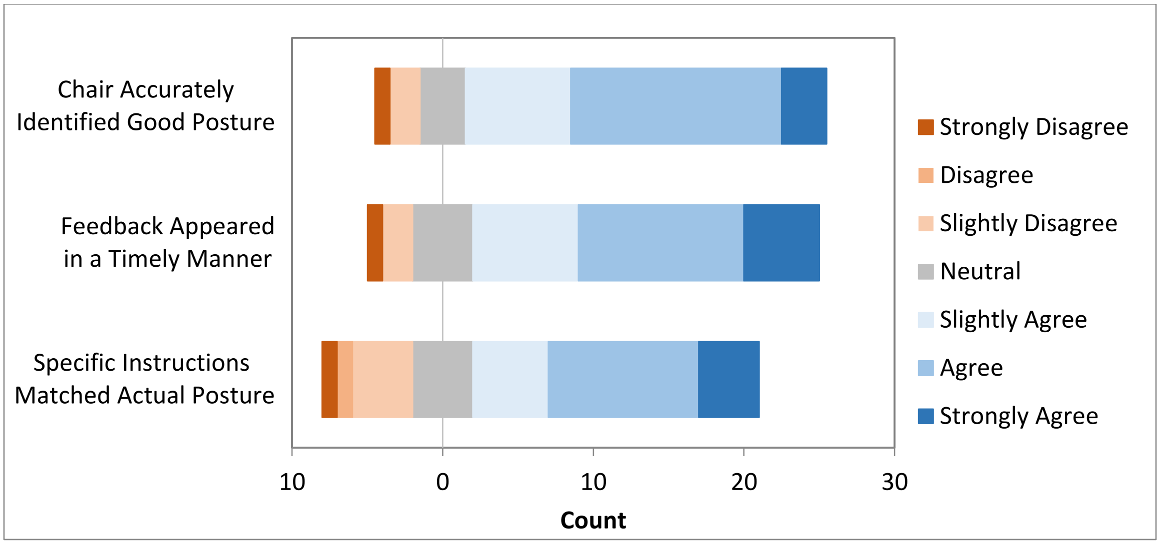

Figure 21. Participant agreement with statements regarding the accuracy of the feedback. 30 participants in total.

5.3.3 Participant Preference on Type of Feedback Pop-Up

One of the main aims of the questionnaire was to determine the participants' preference for one pop-up over the other. Participants were asked Likert questions about whether they found either of the pop-ups helpful. They were also asked outright which pop-up they preferred, Generic or Specific, and then asked to explain their choice in a long-answer format. Participants overwhelmingly preferred the Specific feedback, with only 5 out of 30 participants preferring the Generic feedback. One participant chose "Both" even though that was not a stated option. This overwhelming preference for the Specific feedback is supported by the participants' agreement with the statements "The feedback with specific instructions helped me improve my posture" and "The feedback without specific instructions helped me improve my posture". Comparing the responses to these questions with the non-parametric Wilcoxon signed-rank test, a significantly larger number of participants agreed more strongly with the statement that the Specific feedback helped improve their posture (24 participants) than agreed more strongly with the statement that the Generic feedback helped improve their posture (2 participants), T = 38.5, p = .000. Notably, 4 participants agreed equally that each of the feedback pop-ups helped improve their posture. Clearly participant preference for the Specific feedback was strong. The group that preferred the Generic feedback is large enough, however, that the ability to choose their preferred feedback type might be an appealing feature in a commercial system.

Participants were able to communicate the reasons for their preference of one feedback type over the other in a long-answer question on the questionnaire. A strong theme emerged from those who preferred the Specific feedback: they felt that the extra information presented in that feedback pop-up was necessary to enable them to return to the reference posture. Several participants expressed that they had trouble being aware of their body, or did not know exactly how they should correct their posture, and that they appreciated that the feedback helped them "avoid the possibility that I might change my posture to another poor one". A few enjoyed the challenge of getting all of the "dots" on the Specific feedback window from red to grey, which increased their motivation to match the reference posture. Those who preferred the Generic feedback mostly felt that they already knew what a good posture was, and so did not need the extra information. One participant found the Specific feedback too distracting. Several participants wanted a more customizable system, for instance making the Specific feedback an optional window that they could look at only when they chose to, or having Specific feedback as the first few pop-ups and then switching to Generic feedback once they "get used to the chair and the proper way to sit". The participant who preferred "Both" feedback pop-ups said they would prefer the Generic pop-up for daily use, but that the Specific instructions would be helpful “after the software has had time to create an accurate profile of the user’s posture”. This suggestion that the system would learn about the user’s posture habits over time and thus improve the accuracy of its feedback would be an interesting improvement to the current system that could potentially make it even more effective.

5.3.4 Possible Improvements to Feedback

The feedback pop-ups in the PostureChair system were designed to be very simple to determine the basic effectiveness of digital feedback. However, participants were asked on the questionnaire to suggest any additional or different information they would like to see on the feedback pop-ups to provide direction for future work. Several participants used this opportunity to suggest changes to the feedback pop-ups, usually to make them less disruptive. Suggestions included having the feedback in the system tray rather than as pop-ups, having pop-ups that did not require dismissal but dismissed themselves either after a time-lapse or because the user had achieved the correct posture, and pop-ups that stuck to the edge of the screen rather than appearing in the middle. One participant mentioned that they did not find the pop-ups themselves disruptive, but that the delay between when the pop-up appeared and when it could be dismissed was too long. Since the pop-ups were designed to be as obvious as possible, it is not surprising that more subtle feedback was suggested. In a commercial system that is used continuously, it is also likely that the feedback pop-ups in this system would become more irritating than they did in a 2-hour user study. An extended study in an office setting would be beneficial to determine the longer-term, real-world effectiveness of our system and to examine more subtle digital feedback.

Many participants suggested that the feedback should give clear goals to reach to provide motivation for improving posture. They also suggested additional information should be displayed, including indicators of how long the user had been sitting overall, how much time they had spent in good or bad posture that day or week, and the number of pop-ups they had seen that day or week. These numbers would allow users to track their behaviour over weeks or months. A few participants were interested in having the system give positive feedback, either as words of praise or other rewards based on actual improvement in their posture. Positive feedback would add a new method of persuasion, Conditioning, to the feedback in our system. Long term scores and rewards are hallmarks of gamification, which two participants actually suggested explicitly. Gamification is a common technique in persuasive technology and involves turning the improvement of the behaviour into a game in which users can "win" rewards, often simply points or badges, by achieving certain goals.

Several participants felt that the binary indicators on the Specific feedback were too simple. A few suggested more variable feedback that would place the poor posture on a spectrum of "slightly poor" to "very poor". The Specific feedback red and grey circles also did not indicate directionality, and several participants expressed frustration that they could not determine whether to put more or less pressure on particular sensors. It was assumed during development that participants would determine this through experimentation with the system, but a more nuanced feedback would likely need to include clearer indicators. All of these suggestions are good candidates for a future version of our system, although any alterations would have to be studied to ensure the feedback is still effective at modifying users' sitting posture.

5.4 Incidental Results

The results discussed above covered the main goal of this research, to determine the effectiveness of digital feedback at improving sitting posture, and the most important corollary findings of participants' evaluations of the Generic and Specific feedback pop-ups. However, the results from our user study can also be extended to two related areas: sitting breaks and overall posture knowledge. As described in Section 2.3.2, sitting breaks are an important part of healthy computer use and thus are related to good sitting posture. In addition, overall knowledge of good posture and ergonomics could influence a participant's habitual posture, and so could affect how well they were able to maintain and return to the reference posture set in the user study.

5.4.1 Sitting Breaks

The user study format followed the sitting guidelines of Rutten et al. [49]: 25 minutes of sitting followed by 5 minutes of standing. This was done to give participants a rest after long periods of sitting, to prevent fatigue that could affect sitting posture in later sessions. It also reflects the status of sitting breaks as a component of healthy computer habits. However, participants were not forced to stand during the 5 minute break period since sitting breaks were not directly under study. Instead, participants were told that they could stand whenever they wished, but that the breaks were there specifically to give them a chance to stand up and stretch if they wanted to.

The data from the 30 participants in the {{Total User Study} – {Double Run #2}} set was examined for periods of time when no force was applied to any of the sensors. If a period of zero force lasted longer than 10 seconds it was considered as a period of standing. Periods of less than 10 seconds were assumed to be caused by the participant shifting briefly in the chair or by periods of standing that were too short to have an effect on the participant's body. 3 participants stood during one of their user study sessions: 1 during the Generic session and 2 during the Specific session. These periods of standing were usually less than a minute or two. Overall, participants remained seated throughout the testing sessions, likely because of their desire to please the researcher and their focus on their own work. By contrast, nearly two thirds of the participants stood up during at least one of the two break periods: 19 during the first break and 18 during the second break. 15 participants stood during both breaks. Participants who stood during a break averaged about 2 minutes 46 seconds of standing out of the 5 minute break.

As stated above one reason breaks were included in the user study was to prevent participants from experiencing fatigue from prolonged sitting, which could potentially affect their posture in later sessions. No significant correlation was found between whether the participant stood during a break and their posture in the subsequent testing session (i.e. the first break compared to Generic session posture, and the second break compared to Specific session posture). However these sessions involved feedback pop-ups specifically designed to affect the participants' posture, so it is likely that any effect of standing during the breaks was hidden by the effect of the feedback.

Several participants indicated that they appreciated the prompt to take a break, both through the long answer questions on the questionnaire (Appendix A) and during informal talk with the researcher. One suggested a "quiet beep" to get their attention to start the break. Participants were asked what additional information could be displayed on the feedback and several suggested information related to sitting breaks, such as how long they had been sitting and instructions on how to stretch during the breaks. This suggests that some participants saw posture and sitting times as two aspects of maintaining healthy computer habits that could both be addressed by one system. One participant was actually more interested in the feedback on how long they had been sitting than on the posture feedback. This participant preferred the Generic feedback over the Specific feedback, and indicated that they did not feel they needed advice on good posture, which suggests they may not be part of the true target audience for a system like the PostureChair. Another participant suggested that if they were receiving many feedback pop-ups over a short period of time, the system should suggest that they take a break, apparently assuming that their poor posture would be related to fatigue.

In general it appears that participants appreciated the reminder to take a break and stretch, and many were aware that sitting for long periods had a negative effect on their posture and overall comfort. Whether this type of system would be effective at persuading users to take breaks from sitting without the encouragement of the researcher is an important area for future work.

5.4.2 Posture Knowledge

There were four questions in the questionnaire (Appendix A) related to the participants' prior knowledge of ergonomics and sitting posture. Participants were asked how much they agreed with the following statements:

- "I know what an ergonomic sitting posture looks like"
- "I try to sit ergonomically when I'm working"
- "I am aware of the consequences of poor posture"
- "I wish I had more information on how to sit ergonomically at work"

Figure 22 shows the responses to these questions for the 30 participants in the {{Total User Study} – {Double Run #2}} set. Participants were evenly split between those who try to sit ergonomically while working and those who do not. Interestingly, more participants agreed that they know what an ergonomic sitting posture looks like than try to sit ergonomically, suggesting that education on good posture is not enough to ensure computer users will actually try to implement that knowledge. However, this result could be partly due to the use of the term "ergonomic", which may have been less familiar to some of the participants than the simpler, but less precise, term of "good posture". Only 2 participants felt that they were not aware of the consequences of bad posture indicating that being aware of the consequences of poor posture is not enough to encourage computer users to try to maintain good posture. Finally, 24 participants agreed to some extent that they wanted more information on how to sit ergonomically in the workplace, despite many of them already knowing the basics of good posture. From conversations with the participants it appears that general ergonomics advice can be too vague and sometimes does not match what they find to be comfortable. Many participants expressed interest in having an expert, either human or technology, come in to their workspace and identify the specific issues that are present there.

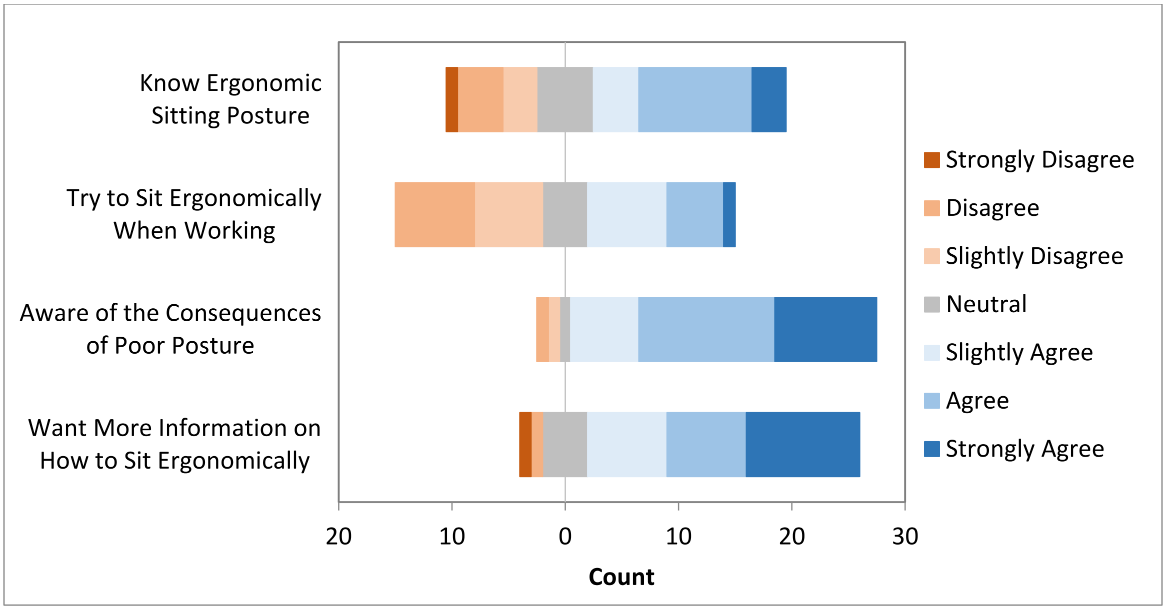

Figure 22. Participant agreement with statements regarding their general knowledge of ergonomics. 30 participants in total.

Participants who agreed more strongly that they know what an ergonomic sitting posture looks like were significantly more likely to agree that they try to sit ergonomically, $r_s = .444$, $p = .014$. This is not surprising since trying to sit ergonomically would be somewhat futile without some knowledge of the posture to aim for. None of the other questions on ergonomics were significantly correlated with each other.

At the beginning of the study participants were asked to sit in an ergonomically good posture to set the reference posture for the PostureChair system. They were given instructions on how to adjust their body into this good posture by the researcher. As shown in Figure 22 many participants agreed that they wished they had more information on how to sit ergonomically. The responses to this question were strongly correlated with the participants' actual posture in the Baseline session, $r_s = .485$, $p = .007$, Generic session, $r_s = .550$, $p = .002$, and Specific session, $r_s = .580$, $p = .009$. In other words, participants who wished they had more information on how to achieve an ergonomic posture spent more time in bad posture regardless of whether they were receiving feedback or not. This suggests that participants that have bad posture generally are aware of it and want the information that will help them improve. While this supports the idea that there are computer users who would be receptive to a system that would help them improve their posture, it also suggests that the information from the PostureChair system was not sufficiently specific. The information provided by the Specific feedback was displayed in abstract terms, and several participants indicated on the questionnaire that they would prefer feedback that showed a representation of their posture in human form instead. One participant suggested a small human "character" that could demonstrate the appropriate movement required to achieve a good posture. Another thought specific instructions in text such as "shift your shoulders left" would be more helpful. The graphic representation of posture in the Specific feedback was a very simple representation of the participant's current posture; these results suggest that more nuanced feedback communicated in terms of the human body could be more appealing and helpful to users.

5.4.3 Perceptions of the PostureChair system

The remaining questions in the questionnaire (Appendix A) address how the participants felt about the system overall. Question 8 ("Before using the chair, I thought my posture was good") and Question 9 ("After using the chair, I think my posture is good") were intended to examine the effect of the PostureChair system on participants' impressions of their typical posture. However, many participants expressed confusion over the meaning of Question 9 in particular, so these two questions will be excluded from analysis. All 30 participants agreed to some extent that they would use the PostureChair system or something similar if it was available in their workplace, with the exception of one participant who was neutral. 24 the participants also agreed to some extent that they expect to pay more attention to their posture after using the PostureChair system. This highlights that the PostureChair system brings awareness to the actual posture users are adopting moment-to-moment. This is far more useful to computer users than the abstract, generalized information provided in ergonomics publications. These encouraging results are subject to an accidental sampling bias since the recruitment methods used for this study (described in Section 4.2) made it more likely that participants would be interested in posture as a topic and in improving their own sitting posture. However, the results of these questions, together with informal discussion between the participants and the researcher, suggest that the PostureChair system addresses a need that is not currently being met in the workplace.

Chapter 6 Conclusions

6.1 Thesis Findings

There have been many attempts at posture detection and modification systems both in academic literature and commercially. These systems have great potential but limited work has been done to determine their effectiveness. We have developed the posture detection and modification system called the PostureChair. This system detects posture through an office chair augmented with force sensors and attempts to persuade users to improve their posture through digital feedback. The posture detection method used in this system was influenced by the most predominant method in the literature. Digital feedback was chosen as the most familiar feedback modality for users and because there are few works examining its effectiveness in the literature. Two different feedback types were developed to determine the amount of information necessary to improve a user's sitting posture and to create an appealing experience.

A within-subjects user study showed that the PostureChair system is effective at modifying a user's sitting posture under both the Generic feedback and more detailed Specific feedback conditions. Contextual, digital feedback proved to be an effective means to persuade users to adopt a more ergonomic posture. The graphical posture representation in the Specific feedback was not strictly necessary to modify users' posture. However, some users did experience greater posture improvement with the more detailed feedback and users preferred the Specific feedback overall. The appearance of the digital feedback windows was perceived to be accurate by the participants and to cause moderately low disruption to workflow. The overall perception of the PostureChair system proved to be positive and the system improved users' awareness of their sitting posture. Though general information on ergonomics is widely available, there is a lack of tailoring of that information to users' specific workplace environments and requirements. The PostureChair system has been shown to be one possible component of the solution to this problem.

Although the PostureChair system was designed with 8 sensors, a code error that was discovered partway through the completion of the user study created an alternative 4-sensor setup. However, there proved to be no significant different in the user experience with these two setups. The 4-sensor setup appeared to be sufficient for binary posture classification and for providing simple feedback that created effective posture improvement, though this requires a more rigorous, purposeful study to be confirmed. No examination of the 4-sensor setup with the more detailed Specific feedback was carried out.

6.2 Limitations and Future Work

As with all persuasive technologies, the effectiveness of the digital feedback in the PostureChair system is limited by the user's willingness to respond. Users can choose to ignore the feedback windows entirely and refuse to change their posture. More persuasive methods can be added to the current system, however this limitation is impossible to fully mitigate without resorting to coercion.

While all possible efforts were made to provide an adjustable desk and office chair for the participants of the user study, the available setup did have limitations. Notably, the keyboard and mouse height were not adjustable relative to the height of the desk. This meant some shorter participants had to compromise on the chair height and several could not achieve the ideal ergonomic posture. There were also discrepancies between the lumbar curve of the chair back and some participants' spines. These limitations reflect the restrictions found in real-world office environments and may have made it more difficult for some participants to achieve a true ergonomic posture. However, the PostureChair's reference posture was set to the limit of what each participant could achieve comfortably. An improved setup would have included a more adjustable desk and chair, as well as multiple chair sizes. This would increase the effectiveness of the system and move the reference posture closer to the ergonomic goal.

Currently, computer users rely on the publications provided by their employer, or more usefully a visit from a specialist, to set up their workstation as ergonomically as possible. The PostureChair system has no ability to detect an improper setup of either desk or chair. A posture detection system that could give feedback on how to change the workstation setup in order to improve the user's chances of achieving an ergonomic posture would be extremely useful. This is not something that is currently found in the literature and may be beyond the current abilities of current detection methods, however we believe it is an important goal for future work.

The reference posture in the PostureChair system is based on a single recording of the sensor data. This increases the chances that the reference posture data may be inaccurate based on a slight movement by the user at the moment of recording. A future system should store the reference posture as an average, taken over a few minutes, to reduce this risk. This would also make it simpler for the user to record their own reference posture, outside of a supervised study.

The classification algorithm in the PostureChair system uses a non-weighted calculation that treats the sensors as interchangeable. This was done as a necessary simplification based on the lack of prior work on possible sensor weightings, however it does not reflect the reality of the variation in pressure on the different sensors. Future work with the aim of improving posture recognition and giving more contextual feedback should examine the validity of different sensor weightings, and the resulting posture recognition rates.

Based on the results of this thesis, and the feedback provided by participants, an all-encompassing posture detection and modification system would include persuasion to increase the number of sitting breaks and display long term data to the user on both posture and break times. Long term goals that the user could work towards would also be a motivating feature. The system would tell the user exactly how to adjust their body to achieve an ergonomic posture, and would detect if the user had set up their chair or desk incorrectly. Additionally, such a system might allow for a range of good postures that users could cycle through during long term use. The PostureChair system is only part of this ideal, but it has already proven to be effective at improving users' sitting posture. This is a promising result and suggests that the all-encompassing system described above would be effective, and would be welcomed by computer users in many workplaces.

Appendices

Appendix A User Study Questionnaire

The following is the questionnaire given to all participants at the end of the user study:

For questions 1-15, please indicate the degree to which you agree with each phrase.

Circle one only.

1) “The feedback popups were **disruptive** to my workflow”

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

2) “The chair **accurately identified** when I had good posture”

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

3) “The feedback popups appeared in a **timely manner**”

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

4) “The feedback popups appeared **too often**”

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

5) “The feedback **with** specific instructions helped me improve my posture”

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

6) “The feedback **without** specific instructions helped me improve my posture”

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

7) “The specific instructions on how to change my posture **matched** my actual posture at the time”

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

8) “**Before** using the chair, I thought my posture was good”

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

9) “**After** using the chair, I think my posture is good”

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

10) “I will likely pay **more attention** to my posture after using this chair”

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

11) “**I would use** a chair like this if it was made available in my workplace”

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

12) "I **know** what an ergonomic sitting posture looks like"

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

13) "I **try to sit ergonomically** when I'm working"

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

14) "I am **aware of the consequences** of poor posture"

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

15) "**I wish I had more information** on how to sit ergonomically at work"

Strongly Disagree	Disagree	Slightly Disagree	Neutral	Slightly Agree	Agree	Strongly Agree

For questions 16 & 17, please answer in the space provided. If you require more space, the researcher can provide you with additional paper.

16) Which do you think you would you **prefer** for **daily use**? Please circle one and explain why you would prefer that type of feedback:

Feedback *with* specific instructions   Feedback *without* specific instructions

17) What **other information**, if any, would you like to see on the feedback popup?

Demographic Information

18) Please select your gender:

- ☐ Female
- ☐ Male
- ☐ Prefer Not to Answer

19) My age is (check one):

- ☐ 18-29
- ☐ 30-39
- ☐ 40-49
- ☐ 50-59
- ☐ 60-69
- ☐ 70+
- ☐ Prefer Not to Answer

20) My height is (check one):

- ☐ Less than or equal to **4’10”** (145cm)
- ☐ Greater than **4’10”** (145cm) but less than or equal to **5’4”** (160cm)
- ☐ Greater than **5’4”** (160cm) but less than or equal to **5’9”** (172cm)
- ☐ Greater than **5’9”** (172cm) but less than or equal to **6’2”** (185cm)
- ☐ Greater than **6’2”** (185cm)
- ☐ Prefer Not to Answer

21) My weight is (check one):

- ☐ **119 lbs** (53kg) or less
- ☐ **120 lbs** (54kg) to **139 lbs** (63kg)
- ☐ **140 lbs** (64kg) to **159 lbs** (72kg)
- ☐ **160 lbs** (73kg) to **179 lbs** (81kg)
- ☐ **180 lbs** (82kg) to **199 lbs** (90kg)
- ☐ **200 lbs** (91kg) to **219 lbs** (99kg)
- ☐ **220 lbs** (100kg) to **239 lbs** (108kg)
- ☐ **240 lbs** (109kg) to **259 lbs** (117kg)
- ☐ **260 lbs** (118kg) or more
- ☐ Prefer Not to Answer

Appendix B User Study Recruitment Documents

B.1 Poster Call for Participants

This poster was distributed online and posted on bulletin boards at Carleton University to recruit participants to the user study.

Carleton
UNIVERSITY
Canada's Capital University

Participate in a Study on Posture in Office Chairs!

Title: PostureChair: A Real-Time, As-Needed Feedback System for Improving the Sitting Posture of Office Workers

We are currently looking for volunteers to participate in an academic study on the effectiveness of visual feedback on posture when sitting in a "smart" office chair.

To Participate in this Study, you must:

- ✓ Have ~2 hrs of work you can do remotely OR on your own Windows laptop you can to bring to the study
- ✓ Be comfortable in the English language
- ✓ Be able to sit in an office chair without pain or discomfort

The study:
This is a 2-hour study. You will be asked to remain seated in the office chair, with the exception of defined break periods, for the entire 2 hours session. You will be allowed to do your own work on a networked Carleton computer or your own laptop, with minor interruptions.

The ethics protocol for this project has been reviewed and cleared by the Carleton University Research Ethics Board 613-520-2517 or ethics@carleton.ca

To participate please email jessica.speir@carleton.ca with the subject "research study" indicating availability.

Jessica.speir@carleton.ca
Subject: research study

Jessica.speir@carleton.ca
Subject: research study

Jessica.speir@carleton.ca
Subject: research study

Jessica.speir@carleton.ca
Subject: research study

Jessica.speir@carleton.ca
Subject: research study

Jessica.speir@carleton.ca
Subject: research study

Jessica.speir@carleton.ca
Subject: research study

Jessica.speir@carleton.ca
Subject: research study

Jessica.speir@carleton.ca
Subject: research study

Jessica.speir@carleton.ca
Subject: research study

Jessica.speir@carleton.ca
Subject: research study

B.2 Online and Email Call for Participants

This text was distributed online through social media and email mailing lists to recruit participants to the user study.

Call for Participants!

We are currently looking for volunteers to participate in an academic study on the *effectiveness of visual feedback on posture* when sitting in a "smart" office chair. The project is titled: "PostureChair: A Real-Time, As-Needed Feedback System for Improving the Sitting Posture of Office Workers"

To participate in this study you must have about 2 hours of your own work you can EITHER do remotely from a networked computer in one of our labs OR on your own Windows laptop you can bring to the lab. Participants must be comfortable in the English language and must be able to sit in an office chair without pain or discomfort for the entire study session.

This is a 2-hour study. You will be asked to remain seated in a smart office chair, with the exception of defined break periods, for the entire session. You will be allowed to do your own work with minor interruptions.

If you are interested, please email Jessica Speir at jessica.speir@carleton.ca for more details on participating.

The ethics protocol for this research has been reviewed and approved by the Carleton University Research Ethics Board, 613-520-2517 or ethics@carleton.ca.

Appendix C User Study Consent Form

This form was given to participants at the beginning of the user study to show approval from the Ethics department and to collect participant consent.

Carleton
UNIVERSITY
Canada's Capital University

Consent Form

Title: PostureChair: A Real-Time, As-Needed Feedback System for Improving the Sitting Posture of Office Workers

Date of ethics clearance: April 8, 2015

Ethics Clearance for the Collection of Data Expires: May 31, 2016

I ____, choose to participate in a study on the effects of visual feedback on sitting posture. This study aims to compare different kinds of visual feedback and determine their effect on the sitting posture of office workers. **The researcher for this study is Jessica Speir in the Human-Computer Interaction department.** She is working under the supervision of Dr. Anthony Whitehead in the School of Information Technology.

This study involves one 90-minute period with a questionnaire at the end. The session will be broken into three consecutive 25-minute sessions, with 5 minute breaks in between. One session will involve feedback on the quality of your posture, one will involve feedback on how to improve your posture, and the third will have no feedback. After the sessions are complete, you will be asked to complete a 10-minute questionnaire regarding your experience in the office chair and the feedback that you received.

As this project will not ask you for any personal information, there is minimal risk to you, the participant. All information provided by you will be confidential but not anonymous; your session results and questionnaire responses will be linked to your name but they will not be shared. You may request that certain questionnaire responses not be included in the final project.

This project requires you to follow instructions on your sitting posture, so there is minimal risk of back strain. To further minimize this risk, we ask that you exercise caution: ignore any instructions that you think might cause strain and to stop if you experience any discomfort.

You have the right to end your participation in the study at any time, for any reason, up until July 1st, 2015. You can withdraw by phoning or emailing the researcher or the research supervisor. If you withdraw from the study, all information you have provided will be immediately destroyed.

All research data, including logs and any notes will be password-protected. Any hard copies of data (including any handwritten notes or USB keys) will be kept in a locked cabinet at Carleton University. Research data will only be accessible by the researchers and the research supervisor.

Once the project is completed, all research data will be kept for five years and potentially

Page 1 of 2

This document has been printed on both sides of a single sheet of paper.
Please retain a copy of this document for your records.

used for other research projects on this same topic. At the end of five years, all research data will be securely destroyed. (Electronic data will be erased and hard copies will be shredded.)

If you would like a copy of the finished research project, you are invited to contact the researchers to request an electronic copy which will be provided to you.

This project was reviewed by the Carleton University Research Ethics Board, which provided clearance to carry out the research. Should you have questions or concerns related to your involvement in this research, please contact:

REB contact information:
Professor Louise Heslop, Chair
Professor Andy Adler, Vice-Chair
Research Ethics Board
Carleton University
1325 Dunton Tower
1125 Colonel By Drive
Ottawa, ON K1S 5B6
Tel: 613-520-2517
ethics@carleton.ca

Researcher contact information:
Jessica Speir, Masters Candidate
Carleton University
School of Information Technology
[redacted]
Email: jessica.speir@carleton.ca

Supervisor contact information:
Dr. Anthony Whitehead
Carleton University
School of Information Technology
[redacted]
Email: anthony.whitehead@carleton.ca

Do you agree to have your interactions with the office chair logged: __Yes __No

______________________ ______________
Signature of participant Date

______________________ ______________
Signature of researcher Date

This document has been printed on both sides of a single sheet of paper.
Please retain a copy of this document for your records.

Appendix D Summary of Responses to Likert Questions

The answers to these questions were on a 7-point Likert scale from "Strongly Disagree", coded as 1, to "Strongly Agree", coded as 7. Significant results shown in bold.

	Descriptive			K-S Test	
	Mean	*Std Dev*	*Mdn*	*D*(39)	*p*
"The feedback pop-ups were disruptive to my workflow"	3.23	1.81	3	0.235**	**.000**
"The chair accurately identified when I had good posture"	5.30	1.32	6	0.269**	**.000**
"The feedback popups appeared in a timely manner"	5.20	1.39	6	0.226**	**.000**
"The feedback pop-ups appeared too often"	3.13	1.55	2.5	0.268**	**.000**
"The feedback with specific instructions helped me improve my posture"	5.60	1.40	6	0.379**	**.000**
"The feedback without specific instructions helped me improve my posture"	4.20	1.47	5	0.240**	**.000**
"The specific instructions on how to change my posture matched my actual posture at the time"	4.80	1.81	8	0.213**	**.001**
"Before using the chair, I thought my posture was good"	3.40	1.71	3	0.193**	**.006**
"After using the chair, I think my posture is good"	4.57	1.52	5	0.212**	**.001**

	Descriptive			K-S Test	
	Mean	*Std Dev*	*Mdn*	*D*(39)	*p*
"I will likely pay more attention to my posture after using this chair"	5.40	1.59	6	0.247**	**.000**
"I would use a chair like this if it was made available in my workplace"	6.10	0.92	6	0.269**	**.000**
"I know what an ergonomic sitting posture looks like"	4.63	1.73	5	0.218**	**.001**
"I try to sit ergonomically when I'm working"	4.00	1.55	4	0.173*	**.022**
"I am aware of the consequences of poor posture"	5.80	1.22	6	0.265**	**.000**
"I wish I had more information on how to sit ergonomically at work"	5.53	1.52	6	0.187**	**.009**

References

[1] Algonquin College Office Ergonomics Program: 2007. *http://www.algonquincollege.com/safety-security-services/files/2014/06/PROG-OHSADM01-Office-Ergonomics.pdf*.

[2] Amick, B.C., Robertson, M.M., DeRango, K., Bazzani, L., Moore, A., Rooney, T. and Harrist, R. 2003. Effect of office ergonomics intervention on reducing musculoskeletal symptoms. *Spine*. 28, 24 (2003), 2706–2711.

[3] Arduino - ArduinoBoardMega2560: *http://www.arduino.cc/en/Main/arduinoBoardMega2560*. Accessed: 2015-05-04.

[4] Arduino - Software: *http://www.arduino.cc/en/Main/Software*. Accessed: 2014-11-28.

[5] Atlas, S.J. and Deyo, R.A. 2001. Evaluating and Managing Acute Low Back Pain in the Primary Care Setting. *Journal of General Internal Medicine*. 16, 2 (2001), 120–131.

[6] Axia Smart Chair: *http://www.bma-ergonomics.com/en/product/axia-smart-chair/*. Accessed: 2015-01-16.

[7] Berque, D., Burgess, J., Billingsley, A., Johnson, S., Bonebright, T.L. and Wethington, B. 2011. Design and evaluation of persuasive technology to encourage healthier typing behaviors. *Proceedings of the 6th International Conference on Persuasive Technology (PERSUASIVE'11)* (2011).

[8] Claus, A.P., Hides, J.A., Moseley, G.L. and Hodges, P.W. 2009. Is "ideal" sitting posture real?: Measurement of spinal curves in four sitting postures. *Manual Therapy*. 14, 4 (2009), 404–408.

[9] Computer Ergonomics: Workstation Layout and Lighting: 2004. *http://www.labour.gov.on.ca/english/hs/pdf/gl_comp_erg.pdf*. Accessed: 2015-07-10.

[10] Daian, I., van Ruiten, A.M., Visser, A. and Zubic, S. 2007. Sensitive chair: a force sensing chair with multimodal real-time feedback via agent. *Proceedings of the 14th European conference on Cognitive ergonomics (ECCE '07)* (2007), 163–166.

[11] DeRango, K., Amick, B.C. and Robertson, M. 2003. *The Productivity Consequences of Two Ergonomic Interventions*.

[12] Diffrient, N., Tilley, A.R. and Bardagjy, J.C. 1974. *Humanscale 1/2/3: a portfolio of information*. MIT Press.

[13] Van der Doelen, L.H.M., Netten, M.P. and Goossens, R.H.. 2011. Tactile feedback to influence sitting behaviour during office work. *Proceedings of the 43rd Annual Conference of the Nordic Ergonomics Society (NES 2011)* (2011), 380–385.

[14] Dunne, L.E., Walsh, P., Hermann, S., Smyth, B. and Caulfield, B. 2008. Wearable monitoring of seated spinal posture. *IEEE Transactions on Biomedical Circuits and Systems*. 2, 2 (2008), 97–105.

[15] Epstein, R., Colford, S., Epstein, E., Loye, B. and Walsh, M. 2012. The effects of feedback on computer workstation posture habits. *Work*. 41, 1 (2012), 73–79.

[16] Ergonomic Workspace Planner: *http://www.ergotron.com/tabid/305/language/en-us/default.aspx*. Accessed: 2015-07-14.

[17] FitBit: 2015. *https://www.fitbit.com/ca*. Accessed: 2015-06-16.

[18] Fogg, B.J. 2009. Creating Persuasive Technologies : An Eight-Step Design Process. *Proceedings of the 4th International Conference on Persuasive Technology (Persuasive '09)* (2009).

[19] Fogg, B.J. 2002. *Persuasive Technology: Using Computers to Change What We Think and Do (Interactive Technologies)*.

[20] Force Sensitive Resistor - Square - SparkFun Electronics: *https://www.sparkfun.com/products/9376*. Accessed: 2015-09-04.

[21] Goossens, R.H.., Netten, M.P. and van der Doelen, B. 2012. An office chair to influence the sitting behavior of office workers. *Work*. 41, (2012), 2086–2088.

[22] Griffiths, E. and Saponas, T.S. 2014. Health Chair: Implicitly Sensing Heart and Respiratory Rate. *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '14)* (2014), 661–671.

[23] Grondin, D.E., Triano, J.J., Tran, S. and Soave, D. 2013. The effect of a lumbar support pillow on lumbar posture and comfort during a prolonged seated task. *Chiropractic & Manual Therapies*. 21, 21 (2013).

[24] Haller, M., Richter, C., Brandl, P., Gross, S., Schossleitner, G., Schrempf, A., Nii, H., Sugimoto, M. and Masahiko, I. 2011. Finding the Right Way for Interrupting People Improving Their Sitting Posture. *Human-Computer Interaction - INTERACT 2011 Part II*. Springer Berlin Heidelberg. 1–17.

[25] Hamilton, M.T., Healy, G.N., Dunstan, D.W., Zderic, T.W. and Owen, N. 2008. 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.

[26] Hart, S.G. and Staveland, L.E. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. *Advances in psychology*. 52, (1988), 139–183.

[27] Ishimatsu, H. and Ueoka, R. 2014. BITAIKA: development of self posture adjustment system. *Proceedings of the 5th Augmented Human International Conference - AH '14* (2014).

[28] Jawbone UP3: 2015. *https://jawbone.com/store/buy/up3*. Accessed: 2015-06-16.

[29] Kinect for Windows: 2015. *http://www.microsoft.com/en-us/kinectforwindows/*. Accessed: 2015-06-08.

[30] Klasnja, P. and Pratt, W. 2012. Healthcare in the pocket: mapping the space of mobile-phone health interventions. *Journal of biomedical informatics*. 45, 1 (2012), 184–98.

[31] Liang, G., Cao, J., Liu, X. and Han, X. 2014. Cushionware: a practical sitting posture-based interaction system. *CHI '14 Extended Abstracts on Human Factors in Computing Systems (CHI EA '14)* (Apr. 2014), 591–594.

[32] Lumo Back: 2014. *http://www.lumobodytech.com/lumoback/*. Accessed: 2015-06-18.

[33] Lumo Lift: 2014. *http://www.lumobodytech.com/*. Accessed: 2015-01-16.

[34] Marshall, K. 2001. Statistics Canada: Working with computers. *Perspectives on Labour and Income*. 2, 5 (2001), 1–23.

[35] Martins, L., Lucena, R. and Belo, J. 2013. Intelligent Chair Sensor. *Engineering Applications of Neural Networks*. Springer Berlin Heidelberg. 182–191.

[36] MSDN Library Kinect Skeletal Tracking: 2015. *https://msdn.microsoft.com/en-us/library/hh973074.aspx*. Accessed: 2015-07-10.

[37] Mutlu, B., Krause, A., Forlizzi, J., Guestrin, C. and Hodgins, J. 2007. Robust, low-cost, non-intrusive sensing and recognition of seated postures. *Proceedings of the 20th annual ACM symposium on User interface software and technology (UIST '07)* (2007), 149–158.

[38] Netten, M.P., van der Doelen, L.H.M. and Goossens, R.H.M. 2013. Chair Based Measurements of Sitting Behavior a Field Study of Sitting Postures and Sitting Time in Office Work. *Digital Human Modeling and Applications in Health, Safety, Ergonomics, and Risk Management. Human Body Modeling and Ergonomics*. Springer Berlin Heidelberg. 261–268.

[39] Van Niekerk, S.-M., Louw, Q., Vaughan, C., Grimmer-Somers, K. and Schreve, K. 2008. Photographic measurement of upper-body sitting posture of high school students: a reliability and validity study. *BMC musculoskeletal disorders*. 9, (2008), 113.

[40] Obermair, C., Reitberger, W., Meschtscherjakov, A., Lankes, M. and Tscheligi, M. 2008. perFrames: Persuasive picture frames for proper posture. *Persuasive Technology*. Springer Berlin Heidelberg. 128–139.

[41] Oinas-Kukkonen, H. 2010. Behavior change support systems: A research model and agenda. *Proceedings of the 5th International Conference on Persuasive Technology (PERSUASIVE '10)* (2010), 4–14.

[42] OSH Answers - Good Body Position: 2010. *http://www.ccohs.ca/oshanswers/ergonomics/sitting/sitting_position.html*. Accessed: 2015-07-10.

[43] OSH Answers - Working in a Sitting Position: 2010. *http://www.ccohs.ca/oshanswers/ergonomics/sitting/sitting_overview.html*. Accessed: 2015-06-01.

[44] OSHA Computer Workstation eTool: *https://www.osha.gov/SLTC/etools/computerworkstations/positions.html*. Accessed: 2015-07-10.

[45] Owen, N., Bauman, A. and Brown, W. 2009. Too much sitting: a novel and important predictor of chronic disease risk? *British journal of sports medicine*. 43, 2 (2009), 81–83.

[46] Owen, N., Healy, G.N., Matthews, C.E. and Dunstan, D.W. 2010. Too much sitting: the population-health science of sedentary behavior. *Exercise and sport sciences reviews*. 38, 3 (2010), 105–113.

[47] Philips ErgoSensor Monitor: *http://ergosensormonitor.com/lib/docs/241P4LRYES_00.pdf*. Accessed: 2015-08-06.

[48] RSI Guard Complete User Information: 2013. *http://rsiguard.remedyinteractive.com/support/individual-rsiguard-users/rsiguard-complete-user-information/*. Accessed: 2015-06-02.

[49] Rutten, G.M., Savelberg, H.H., Biddle, S.J.H. and Kremers, S.P.J. 2013. Interrupting long periods of sitting: good STUFF. *The International Journal of Behavioral and Nutrition and Physical Activity*. 10, 1 (2013).

[50] Ryan, C.G., Dall, P.M., Granat, M.H. and Grant, P.M. 2011. Sitting patterns at work: objective measurement of adherence to current recommendations. *Ergonomics*. 54, 6 (2011), 531–538.

[51] Saga, T., Munekata, N. and Ono, T. 2014. Daily support robots that move on me. *SIGGRAPH Asia 2014 Emerging Technologies (SA ’14)* (Nov. 2014).

[52] Schrempf, A., Schossleitner, G., Minarik, T., Haller, M., Gross, S. and Kurschl, W. 2011. PostureCare - Towards a novel system for posture monitoring and guidance. *Proceedings of the 18th IFAC World Congress (IFAC ’11)* (2011).

[53] Sonne, T. and Grønbæk, K. 2013. Exploring new potentials in preventing unhealthy computer habits. *CHI ’13 Extended Abstracts on Human Factors in Computing Systems (CHI EA '13)* (2013), 487–492.

[54] Statistics Canada 2011. *General Social Survey - 2010 Overview of the Time Use of Canadians*.

[55] Tan, H.Z., Slivovsky, L.A. and Pentland, A. 2001. A Sensing Chair Using Pressure Distribution Sensors. *IEEE/ASME Transactions on Mechatronics*. 6, 3 (2001), 261–268.

[56] Taylor, B., Birk, M., Mandryk, R.L. and Ivkovic, Z. 2013. Posture training with real-time visual feedback. *CHI ’13 Extended Abstracts on Human Factors in Computing Systems (CHI EA '13)* (2013), 3135–3138.

[57] The Ergonomic Seating Guide: 2008. *http://www.haworth.com/docs/default-source/documents-files-very-task/ergonomic_seating_guide_handbook-pdf-13337.pdf*. Accessed: 2014-07-18.

[58] Wahlström, J. 2005. Ergonomics, musculoskeletal disorders and computer work. *Occupational Medicine*. 55, 3 (2005), 168–176.

[59] WCB Alberta Office Ergonomics: *http://www.wcb.ab.ca/pdfs/public/office_ergo.pdf*. Accessed: 2015-08-06.

[60] Workplace Safety and Insurance Board of Ontario 2013. *By the Numbers: 2013 WSIB Statistical Report - Schedule 1*.

[61] WorkSafeNB Office Ergonomics: 2010. .

[62] Zheng, Y.J. and Morrell, J.B. 2010. A Vibrotactile Feedback Approach to Posture Guidance. *2010 IEEE Haptics Symposium* (2010), 351–358.