applied system innovation

Article

Design, Implementation, and Field Testing of a Privacy-Aware Compliance Tracking System for Bedside Care in Nursing Homes

Wenbing Zhao [1,*], Qing Wu [1], Ann Reinthal [2] and Nian Zhang [3]

1 Department of Electrical Engineering and Computer Science, Cleveland State University, Cleveland, OH 44115, USA; q.wu55@vikes.csuohio.edu
2 School of Health Sciences, Cleveland State University, Cleveland, OH 44115, USA; a.karas@csuohio.edu
3 Department of Electrical and Computer Engineering, University of the District of Columbia, Washington, DC 20008, USA; nzhang@udc.edu
* Correspondence: wenbing@ieee.org; Tel.: +1-216-523-7480

Received: 29 November 2017; Accepted: 18 December 2017; Published: 22 December 2017

Abstract: Lower back musculoskeletal disorders are pervasive in workplaces. In the United States alone, the total cost of such injuries exceed $100 billion a year. The lower-back injury rate in the healthcare sector is one of the highest among all industry sectors. A main risk factor for lower-back injuries is the use of improper body mechanics when doing lifting and pulling activities. In healthcare venues, nursing homes in particular, nursing assistants are on the front line to take care of patients. Even in places where ceiling-mounted lifting equipment is installed, they are still required to handle the patient for bedside care, such as sliding the sling underneath the patient, scooping up the patient, putting on compression socks, etc. To help nursing assistants get into the habit of using proper body mechanics, we designed and implemented a privacy-aware compliance tracking system (PACTS). PACTS can track a nursing assistant for possible violation of proper body mechanics while doing bedside care and provide realtime feedback via a smart wearable device such as a smart watch worn by the nursing assistant. The system was deployed in a local nursing home for an 80-day field study in six rooms with seven participating nursing assistants. The test exposed several issues with the original design of the system. The primary issue is how to balance the privacy requirement and the usability of the system. Over-emphasizing the former would negatively impact the latter. This issue is partially resolved with a leasing mechanism where the system would automatically register a nursing assistant within the lease period once she or he has manually registered with the system.

Keywords: human activity tracking; microsoft kinect; smartwatch; smartphone; gesture recognition; human computer interaction

1. Introduction

Lower back musculoskeletal disorders are pervasive in workplaces. In the United States alone, the total cost of such injuries exceed $100 billion a year [1]. The lower-back injury rate in the healthcare sector is one of the highest among all industry sectors. A main risk factor for lower-back injuries is the use of improper body mechanics when doing lifting and pulling activities [2,3]. In healthcare venues, nursing homes in particular, nursing assistants are on the front line to take care of patients. Their lower-back injury rate is more than seven times as high as the average for all workers (249 per 10,000 in healthcare compared to 34 per 10,000 [4]). Even in places where ceiling-mounted lifting equipment is installed, they are still required to handle the patient for bedside care, such as sliding the sling underneath the patient, scooping up the patient, putting on compression socks, etc. To help nursing assistants get into the habit of using proper body mechanics, we designed and implemented a privacy-aware compliance tracking system (PACTS). PACTS can track a nursing assistant for possible violation of proper body mechanics while doing bedside care and provide realtime feedback via a smart wearable device such as a smart watch worn by the nursing assistant. The system can be deployed into multiple rooms in a nursing home, and can follow a particular nursing assistant when she or he goes to different rooms to take care patients. The system was deployed in a local nursing home for an 80-day field study.

Human motion tracking in PACTS is via the computer vision technology using inexpensive commercial-off-the-shelf Microsoft Kinect sensor. Enabled by the Microsoft Kinect Software Development Kit (SDK), PACTS can track up to two users with one Kinect for Xbox/Windows, and up to four users with one Kinect for Xbox One with full skeleton joint positions [5]. The accuracy of the Kinect sensor for human motion tracking has been well-established [5,6]. However, this does not mean that it is straightforward to design a Kinect-based system for our purpose.

There are two obvious challenges. First, while it is straightforward to use Kinect and its SDK to track human motion, additional mechanisms are required for Kinect-based systems to be used in healthcare venues because of the privacy of the patients must be respected and Kinect-based tracking is promiscuous, i.e., everyone who is in the view of the Kinect sensor might be tracked. Second, a major challenge of using the computer vision technology in nursing homes is to ensure reliable human motion tracking despite possibly severe occlusions caused by furniture such as bed, couch, walker, and wheelchair. Fortunately, in our target venue, only the nursing assistant's lower extremity is often occluded and, hence, the occlusion does not preclude us from detecting back-bending activities because we can use the hip center and the should center joints to estimate the trunk orientation with respect to the floor.

On detection of a non-compliant activity that may increase the risk of lower-back pain, we must also design an effective intervention mechanism that alerts the nursing assistants about the event, but does not interfere with their job activities. During the field testing, we encountered another challenge not expected at the beginning, that is, how to balance the need for privacy protection and the usability of the system. Making the privacy protection stronger will inevitably reduce the usability of the system.

PACTS incorporates several innovative mechanisms to cope with these challenges:

- A privacy-aware mechanism that ensures only the consented users will be tracked even when she or he works in several different rooms during her/his shift, which protects the privacy of patients and other unrelated persons who might come to the view of the Kinect sensor.
- An occlusion-resilient human activity tracking mechanism that works well in crowded rooms in nursing homes.
- A mechanism to deliver personalized and discreet alerts to the user on detection of a non-compliant activity in realtime.
- A lease-based mechanism that strikes a good balance between meeting the privacy requirement and the usability of the system.

2. Background and Related Work

There is a vast literature on human activity tracking using computer vision. For a comprehensive review of the research on human activity tracking using Microsoft Kinect, readers are referred to [5,7]. Human activity recognition is typically accomplished via two approaches: (1) learning based; and (2) rule based. The learning-based approach relies on the existence of pre-collected and pre-labeled data for training or template comparison [8,9]. The main benefit of the learning-based approach is that no expert knowledge is needed to define the activity to be recognized. The downside is the lack of precise information such approach can provide. Quite often, the detection (i.e., the classification) of an activity is insufficient. For example, on the back-ending activity, the degree of bending contains very valuable information that should be recorded and reported. The rule-based approach requires expert knowledge on the activity to be detected, and it also requires test runs to find the right error bounds on the activity detection. The rule-based approach is less computationally intensive and could capture precise information regarding the detected activity, such as the back-flexion angle. For PACTS, we choose to follow a rule-based approach that we developed previously [6,10–12]. It fits better with the requirement for PACTS than a competing rule-based approach [13] because our rule-based framework could dynamically determine the starting and ending frame of an activity, and it can be integrated seamlessly with our alerting mechanism.

Alerts/alarms have been well studied in the field of network/computer security (e.g., intrusion detection) [14–16] and medicine [17–19]. The main research objective on alerts/alarms generation is to establish the boundary for the normal condition. In the area of medicine, because the target is the individual patient's condition, the decision on when to generate an alarm is even more challenging than that for intrusion detection because the boundary conditions for different patients may vary significantly. The alert generation problem in our study roughly belongs to the area of medicine, although in our system an alert is triggered based on how a worker moves instead of based on his/her vital signs. Alert/alarm generation typically follow the following three approaches: (1) rule based; (2) model based; and (3) machine-learning based.

In the rule-based approach, one or more rules are used to define normal or abnormal conditions. In the security area, this approach has been widely used. For example, one of the most common network-based intrusion tolerance tools, Snort, allows its users to define sophisticated rules for alert/alarm generation [20]. In the area of medicine, rules based on threshold values for a patient's vital signs [17–19] or correct movements [10,11] have been used for alarm generation. A common issue with the rule-based approach is that a large number of alerts/alarms may be generated with high rate of false positives for a complex system. To ensure higher accuracy and reasonable alert/alarm frequency, initial alerts/alarms are often aggregated via correlation [15,16].

In the model-based approach, a scientific model is constructed to predict the behavior of the target system (such as a computer system, a network, or a person) based on the understanding of the system. This approach has been used in some areas of medicine [21,22]. The challenges for this approach include accurate understanding of the system, accurate measurement of the system output, and making the right tradeoff between complexity and efficiency [23].

The machine-learning based approach is appropriate when the target system cannot be accurately modeled, such as the behavior of a computer system/network [24], or some physiological processes of a patient [25,26]. By collecting large amount of training data, statistical models can be trained and used to predict the expected output of the target system [25–27]. It is also possible to use active learning [28] to improve the alert/alarm accuracy by actively taking the inputs from the alert/alarm recipient regarding their quality.

In PACTS, we combine the rule-based and model-based approaches for alert generation. A unique characteristic of PACTS is that the target system and the intended recipient for alerts are the same user. The rule-based approach is used to inspect individual frames with respect to the predefined key posture specification, and the finite state machine is used to model the life-cycle of each non-compliant activity. The combining of these two approaches ensures the generation of one alert for each non-compliant activity.

This study is a continuation of our project started in 2015 [29,30]. Initially, the system was designed to work in a single room only. This limitation would prevent the system from being used in practice because a nursing assistant in a nursing home is often assigned to take care of patients in several different rooms. We later extended the original design so that systems deployed in multiple rooms could be federated together to enable the tracking a nursing assistant when she or he goes to different rooms to do her/his work [31]. In this paper, we provide greater technical details on the system design and present a more comprehensive analysis on the data obtained during the field study.

3. System Design

In this section, we outline the system requirement and elaborate our design rationale.

3.1. System Requirement

To achieve the objective of our project in helping nursing assistants to develop the good habit of using proper body mechanics while providing bedside care, the system should have the following properties:

- Privacy protection. The system must track only the consented nursing assistant, not any other person that might appear in front of the Kinect sensor, including the patient in the room. This is essential to protect the privacy of the patients. The US federal regulation requires that an explicit consent is required before a patient can be tracked. Because it is unlikely for a patient to agree to be tracked all the time for any purpose, the system must provide maximum provision that he or she is not tracked.
- Reliable non-compliance activity recognition. The rooms in a nursing home typically are crowded with limited space. For example, a room may consist of several pieces of furniture such as wheelchairs, tables, walkers, etc. The activity recognition algorithm must be robust despite the occlusion caused by these pieces.
- Realtime alert. To have any positive impact to a nursing assistant, the alert should be delivered in realtime in a way that he or she could accept when he or she performed a non-compliant activity.
- Multi-room support. A nursing assistant is typically assigned to several rooms during a shift. To maximize the effectiveness of our system, we must be able to continuously track the nursing assistant regardless which room he or she goes in to do his or her job duties.
- Balance the privacy protection and usability of the system. While protecting the privacy of the patients, the system must strive to be usable by nursing assistants in a way that is too burdensome for them to use it.

3.2. System Design Rationale

3.2.1. Privacy Protection Design

The rationale for our privacy protection mechanism has been given in our previous work [29,30,32]. We provide a brief overview here. To ensure the privacy protection, we devised a registration mechanism where a consenting nursing assistant must make a specific pose while pressing a button on the smartwatch. On pressing the button, the Kinect runtime would attempt to find a particular user in its view that is making the predefined registration posture. If one and only one is found, that user is identified to be the consenting nursing assistant. This registration posture must be a posture that is not normally made by any nursing assistant or the patient. Hence, the system would only register the consenting nursing assistant.

Furthermore, the system tracks a user only via the skeleton data, not the color video frames, depth frames, or infrared frames. Our system guarantees that only the skeleton data that belong to the tracked user are analyzed for possible non-compliance activities. Our system logs only three types of events: (1) the registration of a user when the tracking of that particular user starts; (2) the user has gone out of the view of the Kinect, or the user has instructed Kinect to stop tracking, when the tracking of the user stops; and (3) when a non-compliance activity is detected.

In addition to the registration mechanism, which relies on the synchronization of two modalities, we also considered other approaches to identify a consenting nursing assistant. One alternative approach is the use of a marker where a nursing assistant would place the marker on her or his uniform. However, this approach may require the marker to be visible all the time, which is not practical because a nursing assistant could be in any position relative to the Kinect sensor. Doing so would also significantly increase the complexity of the system because the system would have to do both skeleton tracking and object detection concurrently the entire time. Another alternative could be the use of NFC (near-field communication) sensors. However, because NFC's operating range is only up to 10 cm, we cannot use an NFC sensor to identify which skeleton is the participating nursing assistant.

In summary, the system does not use color or greyscale frames to operate and does not record such frames in anyway. Hence, an administrator or a hacker cannot misuse the system for any purposes that might violate the privacy of the patients, or track a nursing assistant in a way not intended by our design.

3.2.2. Robust Activity Recognition Algorithm Design

All non-compliance activities that a nursing assistant might engage in while performing bedside care share the same characteristic, i.e., the back is bending too much. Hence, we use the vector determined by the hip center joint and the shoulder center joint, and calculate the angle between this vector and the vector perpendicular to the floor plane, which gives the back flexion angle. This algorithm is robust again occlusion to the lower extremity of a nursing assistant; hence, it can be reliably used in a crowded room.

3.2.3. Realtime Alert Design

On detecting a non-compliance activity, the event is delivered immediately to the smartwatch worn by the nursing assistant, which will generate a vibration alert with a companion text display on the watch's screen. The vibration is discreet in that other nursing assistants or the supervisor will not notice. We hypothesize that this form of alert will be more acceptable by users, which is confirmed by the participants in our field study survey.

3.2.4. Multi-Room Support Design

According to our original design, each room would install one Kinect sensor (with its controlling computer), and a dedicated smartphone would be placed in each room that communicates only with the program running in the computer in the room. Our expectation was that the smartwatch would automatically bind with the dedicated smartphone in the room that the nursing assistant would go into, and unbind with the smartphone placed in the room where the nursing assistant just leaves. Unfortunately, this design did not work because some of the rooms are fairly close to each other, which is within the Bluetooth operating range. As the result, the smartwatch does not do automatic hand-off from one smartphone to another.

Hence, we changed the design to the following. Each participating nursing assistant would carry a smartphone in her or his pocket, and wear a smartwatch on her or his wrist. The smartwatch is bound with the smartphone. The smartphone can communicate with the computer program in every room in which PACTS has been deployed. When a nursing assistant registers with the system, the registration request is broadcast to all rooms and the particular system that could see the nursing assistant would process the request and respond.

This broadcast-based mechanism might appear to be inefficient because it is clearly not scalable. We argue that this is the most desirable design for our use cases. One may suggest that the publish-subscribe pattern could be used for better scalability. Indeed, if we were to broadcast a message to hundreds of destinations, the publish-subscribe approach could be the only viable approach. However, we are targeting nursing homes, which typically have 10–20 rooms that need to have PACTS installed. This is because PACTS is designed to track bedside cares where patients cannot take care of themselves. In a nursing home, only a small fraction of patients are in this category. Hence, we do not need scalability. Furthermore, to enable publish-subscribe, a centralized server would be required. To avoid making this centralized serve a single-point of failure, redundant servers would be needed, which would require sophisticated fault tolerance mechanisms [33–38]. Such an approach would significantly increase the cost and complexity of the system, which is neither needed nor desirable.

In fact, broadcast-based protocols are widely used in practice for self-discovery and low-maintenance, such as the Address Resolution Protocol (ARP) and the Dynamic Host Configuration Protocol (DHCP).

3.2.5. Design for Usability while Protecting Patient Privacy

Privacy is a double-edged sword with respect to the usability of a system. For PACTS, without privacy protection for patients, it will not be allowed to use in nursing homes. Furthermore, without privacy protection for the participating nursing assistants (for example, their images will not be recorded in anyway), they will unlikely agree to be tracked. However, privacy may also negatively impact the usability of the system.

At the beginning of the field study, a nursing assistant would have to manually register with our system every time she or he went to a room to provide bedside care. This caused significant frustration with them and many eventually chose not to register with the system. The nature of the patient-care work requires the nursing assistant to move in and out of a room frequently. Even while a nursing assistant was doing bedside care, she or he might be unregistered by our system due to temporary occlusions by furniture or by the patient. This unexpected difficulty forced us to improve the usability of the system while not significantly reduce the privacy protection of patients.

We decided to incorporate a lease-based mechanism into the system. Once a nursing assistant manually registered with a KinectServer, the KinectServer grants the nursing assistant a lease. While the lease is active, the KinectServer would automatically re-register the nursing assistant as long as the KinectServer has detected one and only one user is present in the view of the Kinect server. This mechanism could improve the usability of the system because the nursing assistant would be automatically re-registered if the KinectServer accidentally unregister her/him or when the nursing assistant goes out and in the view of the Kinect sensor. In the mean time, this mechanism does increase the risk of privacy violation because if one nursing assistant leaves the room and a visitor goes into the same room, the latter would be registered as the nursing assistant. However, this scenario occurs very rarely because a visitor typically would be accompanied by a nursing assistant, at least during the initial visit.

The lease time is critical to the success for the lease-based mechanism. If it is too long, it increases the risk of privacy violation. If it is too short, the usability of the system will not be improved significantly. We consulted with the administrators of the nursing home, and concluded that the lease duration should be set to 30 min because the time corresponds to the average duration of a bedside care session. Hence, with the lease-based mechanism, a nursing assistant would only have to manually register with the system once for every session of bedside care. Typically, a patient only needs bedside care during early morning when she or he is about to get up, and at night when she or he is about to go to sleep. Hence, a nursing assistant would only need to manually register with the system once or twice per shift.

As a pre-caution for privacy protection, we added a way for a nursing assistant to manually terminate the lease and stop tracking when she or he knows that another nursing assistant will replace her/him. For the Pebble smartwatch, this can be done by pressing the down button.

Before we conclude this section, we should note that there are other factors that impact the usability of the system. For PACTS, the non-compliance activity detection accuracy also plays a big role in the usability of the system. As evidenced in the field test, the presence of false positives and false negatives negatively impacted the usability of our system, which we will address in future works.

4. System Implementation

4.1. System Overview

Figure 1 illustrates the PACTS system layout when deployed into multiple rooms. This version of PACTS is built on top of our initial version of the system [32,39,40]. To deploy PACTS in a room, one Kinect sensor and one computer are installed. The Kinect sensor is typically positioned close to the ceiling facing the bed where the patient will rest in. At least during the installation and initial configuration stage, the floor of the room should be clearly visible to the Kinect sensor. The computer runs the KinectServer process that takes input (via a USB cable) from the Kinect sensor and runs all the control logic for user registration, human activity monitoring, and alert reporting.

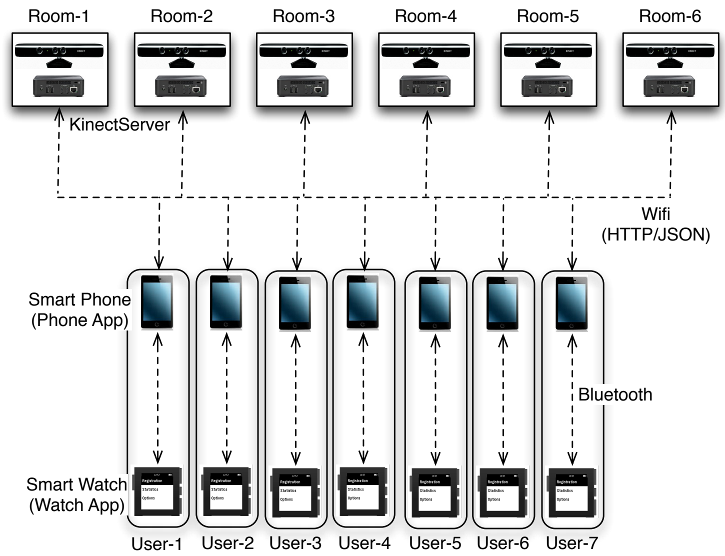

Figure 1. The architecture of the privacy-aware compliance tracking system.

Each user is supposed to wear a smartwatch (we use the Pebble smartwatch) and carry a smartphone (any Android phone or iPhone running the Pebble PhoneApp would work with PACTS). The smartwatch runs a WatchApp we developed and the smartphone runs a corresponding PhoneApp. The smartwatch is connected with the smartphone using Bluetooth. The smartphone connects to all deployed KinectServers via WiFi using JavaScript Object Notation (JSON) over HTTP. The primary role for the PhoneApp is to relay messages between the smartwatch (i.e., the WatchApp) and the KinectServer, although we do incorporate some mechanism to minimize the communication between the smartwatch and the smartphone to reduce the power consumption on both the smartwatch and the smartphone. The energy-saving is more important for the smartwatch because the battery on the smartwatch is quite small compared with that of the smartphone. The PhoneApp is configured to have the IP addresses of all KinectServers and communicate with the KinectServers periodically for the purpose of user registration, user activity detection, and other control needs. The broadcast-based design is necessary because the system has no prior knowledge which room a nursing assistant will go into to work.

The KinectServer runs the main control logic of the system in each room. Hence, it is the most complicated component of PACTS. As shown in Figure 2, the KinectServer takes two forms of inputs:

- Kinect input: The streams of the color and depth frames from connected Kinect sensors.
- Network input: HTTP requests sent by the PhoneApps over the WiFi network.

The Kinect input is used to detect user activities, including the registration gesture and the back-bending activity. The network input contains three types of requests:

- Registration request: It is sent when a user has requested to register with the system. The response for the request indicates whether or not the registration has succeeded. Without the lease mechanism, this request is only initiated by the user via the WatchApp, the registration request will be broadcast to all known KinectServers. When the lease mechanism is enabled, the PhoneApp would periodically sends the registration request to the KinectServer that the user had manually registered with earlier when the user is temporarily unregistered (due to occlusion or due to her/him stepping out of the view of the Kinect sensor).

- Activity status request: It is sent periodically once a user has registered with the system. The response indicates whether or not a back-bending activity has just been detected.
- Stop-tracking request: It is sent when the user wishes the KinectServer to stop tracking temporarily due to privacy concerns, for example, when she is about to give the patient a sponge-bath to make sure that the system absolutely will not track anyone in the room. Once the stop-tracking request is received, the KinectServer also terminates the current lease, if any, for the user.

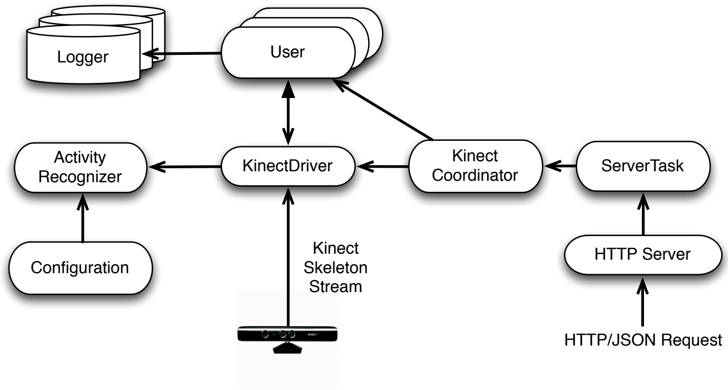

Figure 2. Internal structure of the KinectServer.

Figure 3 illustrates how these types of requests interact with different components of the system. Only the manual registration request is broadcast to all KinectServers. All other types of requests are delivered point-to-point.

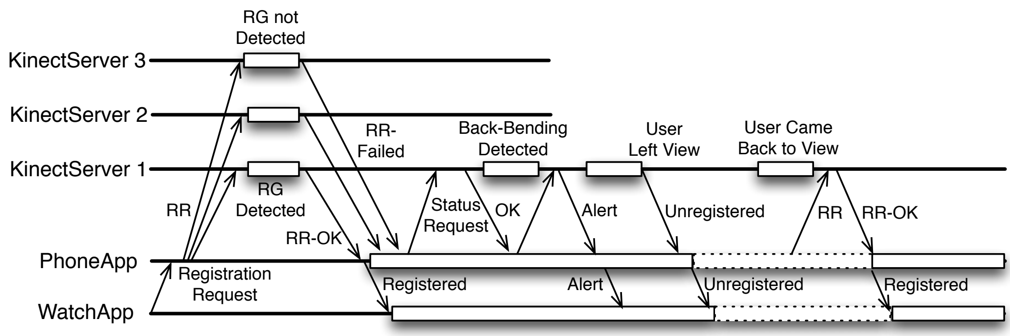

Figure 3. The message flow between different components of the system.

The Kinect server has a number of components:

- Server and ServerTask to handle HTTP communications. The Server component serves as the HTTP listener for the KinectServer. Upon receiving an HTTP request, the Server creates a new ServerTask to handle the message. Both the HTTP request and the response contain a JSON object.
- KinectDriver to handle motion sensing data from the Kinect sensor. In addition, It also handles the user registration request in conjunction with the Activity Recognizer. If the user is properly registered, the user is added to the User object. For each registered user, the KinectDriver tracks her/him for possible non-compliance activities.
- Kinect Coordinator to dispatch the registration and activity requests to a particular KinectDriver and the corresponding User.

During KinectServer initialization, the KinectDriver registers a delegate with the Kinect Sensor so that it can receive the skeleton stream from the Kinect runtime. For each new skeleton frame (i.e., Kinect input), the KinectDriver first makes a copy. If the frame contains a registered user, the KinectDriver invokes the Activity Recognizer to conduct the single-frame level non-compliant activity detection based on what is defined in the Configuration file. The single-frame level activity detection result is passed on to the corresponding User object, which is responsible to carry out the activity-level activity detection.

When a request arrives at the KinectServer from the network (i.e., network input), a ServerTask is created by the HTTP Server object to handle the message. The ServerTask passes the request to the Kinect Coordinator to find the appropriate KinectDriver or User object to handle the request. Then, the corresponding response will be composed and returned to the client.

4.2. User Registration Mechanism

To protect the privacy of patients, a participating nursing assistant is required to register with the system before she or he could be tracked and receive feedback. Before the lease mechanism was introduced, the only way to register with the system is through a manual request by pressing the select button of the Pebble smartwatch while making a predefined gesture. Here we assume that each participating nursing assistant is given a smartwatch exclusively (i.e., not shared with other nursing assistant). The WatchApp generates a unique identifier to represent the user. This identifier is used in all components of the system for the user, including the activity log for the user.

The flow control for user registration at the KinectServer is shown in Figure 4. The registration request arriving at the KinectServer is handled by the Kinect Coordinator. First, the corresponding User object is searched. If one is found (meaning the user has registered with the KinectServer previously), the registration request is passed to the User object for further processing. The User object then invokes the Activity Recognizer to check if the user is making the registration gesture based on what is defined in the Configuration file using the most recent copy of the skeleton frame. If the registration request is for a new user, the request is dispatched to the KinectDriver object, and the KinectDriver subsequently invokes the Activity Recognizer to detect if any of the skeleton contained in the current frame is making the predefined registration gesture. If successful, a User object is created for the user. Finally, a response message is composed indicating whether or not the registration is successful, and is sent back to the PhoneApp and eventually the WatchApp. The registration result is then displayed on the screen of the smartwatch.

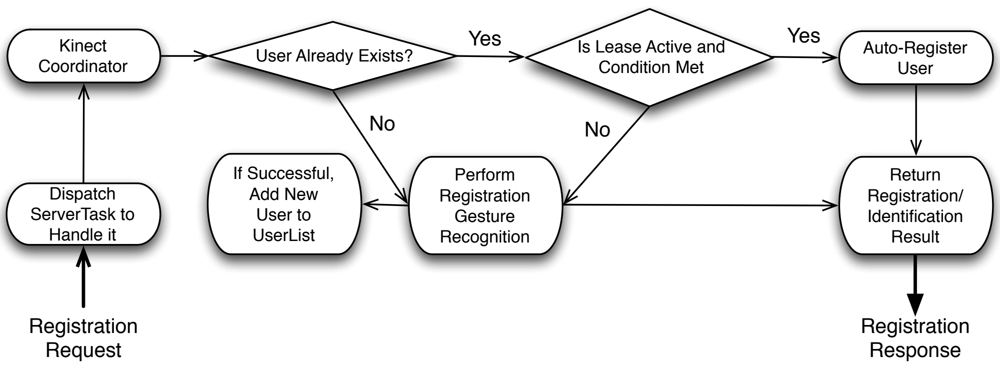

Figure 4. The flow control for user registration at the KinectServer.

The registration mechanism not only protects the privacy of the patient, it enables the system to know who it is tracking and enables continuous tracking of a nursing assistant even when she or he goes to different rooms for bedside (as long as all such rooms are equipped with the KinectServer. In essence, this person-specific tracking is made possible by synchronizing two sensing modalities:

(1) computer vision via the Kinect sensor; and (2) wearable sensing via the smartwatch worn by a user via the predefined registration gesture.

In the current study, we choose to use a registration gesture where the left or right forearm of the nursing assistant pointing straight upward because this gesture allows the her/him to do so using one arm to make the gesture while pressing the select button of the smartwatch using the remaining hand.

When the lease mechanism is enabled, a user is automatically registered before the expiration of the lease provided that the corresponding User object is found in KinectServer and there is one and only one skeleton in the current frame.

4.3. Gesture Recognition and Activity Detection

The Activity Recognizer component in the KinectSever is used to recognize the registration gesture and to detect predefined activities, which are defined in the Configuration file encoded in the eXtensible Markup Language (XML). Both the gesture/activity description and the recognition engine follow the framework that we developed previously [10–12]. In this framework, an activity can be defined in terms of Joint Angle, Joint Distance, and Bone Orientation. The Joint Angle element defines the relative angle between two adjacent body segments represented by three joints (one joint common to both segments, and one each for the other end of the two segments). The Joint Distance element defines the distance between two joints. The Bone Orientation element defines the orientation of a body segment relative to a particular anatomical planes, such as the frontal, sagittal, or transverse plane.

The class hierarchy is shown in Figure 5. The three types of elements are represented by three classes JointAngle, JointDistance, and BoneOrientation. They all implement the abstract class Configuration, which defines a single method called "RecognizeGesture()". To accommodate activities that require the use of complex configurations that use two or more configuration elements, the Gesture class is introduced, which contains the GestureComponents member variable. This member variable is used to fetch and store the list of configurations for the complex gesture/activity.

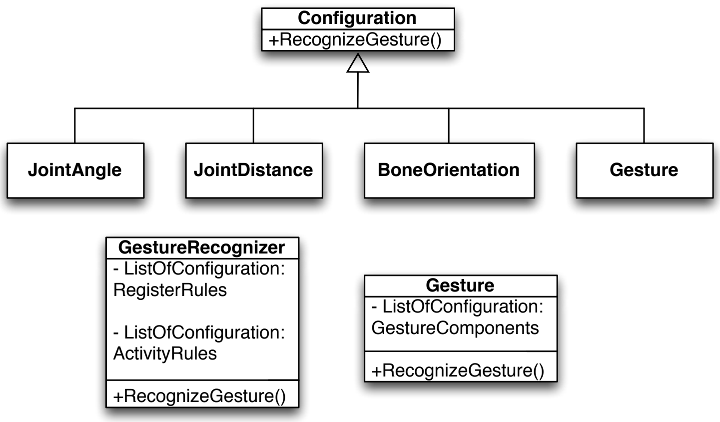

Figure 5. Class diagram for the Activity Recognizer.

The GestureRecognizer class is the front-end class for gesture/activity detection. It has two key member variables: (1) RegisterRules, which contains one or more configurations for the registration gesture; and (2) ActivityRules, which contains the list of configurations for the activities to be detected. The RegisterRules may define multiple alternative gestures such that if any of the gesture is made by the user, the user is considered registered. The GestureRecognizer is created when the system is initiated where the definitions of the gestures and activities are loaded into GestureRecognizer from the configuration file. For the registration gesture, it is detected at the single-frame level. For the back-bending activity, it is detected both on the single-frame level, and the activity level, which we will elaborate further.

In this study, both the registration gesture and the back-bending activity are defined as BoneOrientation elements and they all use the transverse plane as the reference plane, as shown in Listing 1. For the transverse plane, we use the Kinect floor clip plane parameters. This is not only convenient, but also more reliable because it is independent from the user postures. The registration gesture is described in terms of the orientation between the left arm segment from left elbow to the left wrist and the transverse plane, which should be 0 degree relative to the axis that is perpendicular to the transverse plane, with a tolerance of 20 degrees. The back-bending activity is defined in terms of the angle between the vector that is perpendicular to the transverse plane, and the trunk vector from the hip center to the shoulder center. When a standing straight up, the angle should be close to 0. We define the minimum deviation angle of 30 degrees to declare that the user is bending her/his back too much.

Listing 1: Definition for the registration gesture and the back-bending activity used in this study.

```
<RegisterRules>
    <Configuration>
        <Type>BoneOrientation</Type>
        <AlertBody>"Registration"</AlertBody>
        <DownstreamJoint>"WristLeft"</DownstreamJoint>
        <UpstreamJoint>"ElbowLeft"</UpstreamJoint>
        <Plane>"Transverse"</Plane>
        <AlphaAngle>"0"</AlphaAngle>
        <MaxAngleDeviation> "20" </MaxAngleDeviation>
    </Configuration>
</RegisterRules>
<ActivityRules>
    <Configuration>
        <Type>BoneOrientation</Type>
        <AlertBody>"Bend"</AlertBody>
        <DownstreamJoint>"HipCenter"</DownstreamJoint>
        <UpstreamJoint>"ShoulderCenter"</UpstreamJoint>
        <Plane>"Transverse"</Plane>
        <AlphaAngle>"0"</AlphaAngle>
        <MinAngleDeviation> "30" </MinAngleDeviation>
    </Configuration>
</ActivityRules>
```

For both the registration gesture recognition and the non-compliance activity detection, we depend on the floor plane information. As illustrated in Figure 6, we can derive the vector that is perpendicular to the floor plane, V_p. For the registration gesture recognition, we measure the angle A_g formed between the vector from left elbow to the left wrist and V_p. For the back-bending non-compliance activity detection, we measure the back flexion angle A_b formed between the vector from the hip center to the shoulder center and V_p.

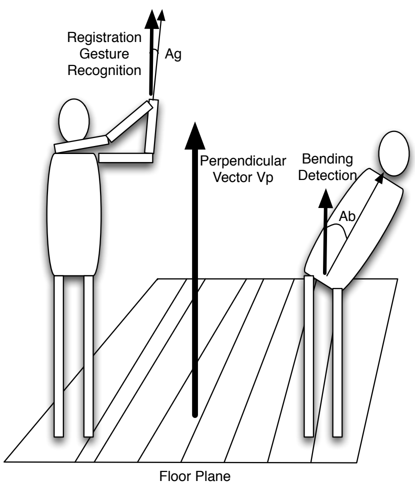

Figure 6. Registration gesture recognition and non-compliant activity detection.

In earlier versions of our system, the floor plane information is obtained from every skeleton frame for gesture/activity recognition. This requires that the floor must always be clearly visible to the Kinect sensor. Our field study showed that this assumption does not hold because the crowded room could prevent Kinect from determining the floor plane. If the floor plane is not visible to Kinect, no one could register with the KinectServer and no non-compliant activity could be detected. This issue is resolved by taking the floor parameters during the initial deployment instead of taking the parameter from each frame. This is possible because the Kinect sensor is fixed at the wall.

4.4. Alerts and Feedback

The purpose of alerts is to remind a consented nursing assistant that she or he is doing or has just done something wrong that could increase the risk of lower back injuries. An alert is generated at the KinectServer when one or more correctness rules are violated, and certain conditions are met. An alert is delivered silently to the worker in the form of a vibration and a short message explaining the nature of the violation (such as back bending) via the WatchApp on the smartwatch. In addition to the alerts, a user is provided with cumulative feedback about her/his performance over time such as by day or by week.

The system can also be configured to log raw activity data in terms of joint positions for each consented worker for institution administrators to review and verify the cumulative worker performance. The log could also be useful for system developers to refine the detection methods for non-compliant activities. To reduce the log space requirement, the system can be configured to log raw data only when a violation is detected.

While our experiments show that a non-compliant activities can be reliably detected by our system, we must follow a carefully crafted method for alert generation and delivery. For example, it would not make any sense if we generate an alert for every frame received that contains a non-compliant posture. For the purpose of our system, we believe that an alert should be generated at the per activity level. As shown in Figure 7, our system generates two types of alerts: active and passive alerts. An active alert is generated on detection of a new non-compliant activity so that the user could terminate the activity as quickly as possible. Hence, an active alert could produce the most benefit to a user by reducing the amount of time the user is in the wrong postures.

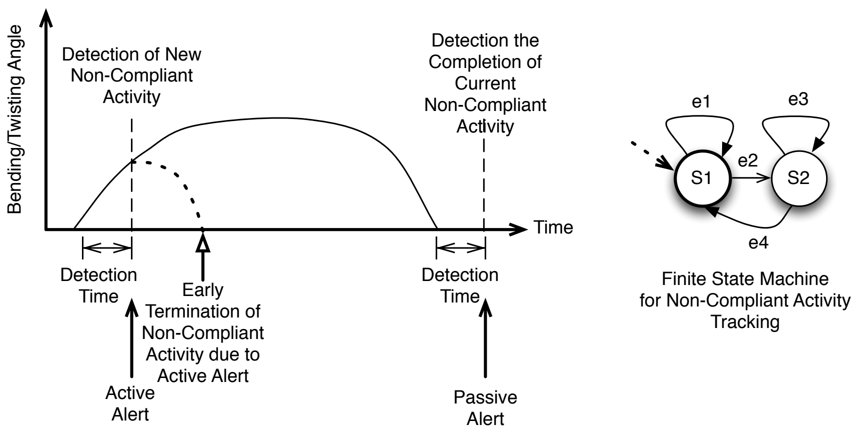

Figure 7. Two types of alerts in our system, and the finite state machine for non-compliant activity tracking.

However, active alerts might not fit all workers and workplaces. First, some nursing assistants may perceive an active alert a distraction and prefer to get the current pulling or lifting job done even if she or he is doing it wrong. Second, the nature of a pulling/lifting job might not be safely terminated once started (such as when a lifting device is used), in this case, an active alert might actually increase safety risks. Hence, our system also offers passive alerts. A passive alert is generated and delivered to the worker as soon as the current non-compliant activity is completed. By delivering a passive alert in between two pulling/lifting jobs, the nursing assistant can take time to review the alert and hopefully remember to do the job correctly the next time.

The types of alerts for each user can be configured by choosing the corresponding alert type via the Options menu in the WatchApp. By default, the active alert is used. Regardless of what alert type is chosen, our system generates a single alert for each non-compliant activity.

Thus far, we have assumed that the starting and ending of a non-compliant activity can be detected reliably and immediately. This task is also critical to ensure that one and only one alert is generated for each non-compliant activity. We decide to use the finite state machine to track the life cycle of a non-compliant activity. A finite state machine for each tracked user has two states as shown on the right side of Figure 7:

- Normal state, denoted as $S1$. It is represented by a by a posture configuration that is within the boundary of the requirement. In our study, it refers to the range when the user is considered to have her/his back straight up.
- Abnormal state, denoted as $S2$. It is represented by a by a posture configuration that is outside the boundary of the requirement. In our study, it refers to the range when the user is considered to have her/his back bent too much.

To be robust against occasional jitters and measurement errors, a transition between $S1$ and $S2$ is made only after a n consecutive frames (in our study, $n = 10$) have all shown the consistent configuration. As a tradeoff for the increased robustness, the detection time for the starting and ending of a non-compliant activity is longer. With $n = 10$, the detection time is 330 ms compared with 33 ms if a single frame is used. This delay is quite acceptable because a typical non-compliant activity lasts several seconds or longer.

5. Field Study

For the field study, we installed PACTS in six rooms at a local nursing home with seven participating nursing assistants. All rooms are single-occupancy rooms where there is only bed and one patient per room. The study lasted 80 days. Because one of the nursing assistants was not assigned to the bedside care duty, only the activities of six nursing assistants were monitored. Furthermore, the patient in one of the rooms did not require bedside care, and another room had several deaths, which made the room unavailable for extended period of time during the field study. Therefore, only the data collected from four rooms are worthy of analysis.

The field study exposed a number of issues with our original design. First, as we mentioned earlier, the manual registration mechanism for privacy protection significantly reduced the usability of the system. Second, the use of a smartphone to relay the communication between the smartwatch and the KinectServer also created usability issues for our system. The smartphone was on the critical path of the registration and alerting process. If the smartphone is in deep sleep and fails to relay the registration request from the smartwatch, the nursing assistant would not be able to register with the KinectServer.

5.1. Analysis of the Logged Data

On Day 35 of the field study, we modified PACTS with the lease-based mechanism after learning the usability issues with our system. The raw data on the total registered duration per room per day in the four rooms are shown in Figure 8. The summary of the data collected is provided in Table 1. As can be seen, only about 888 min (or 13 h) of data are logged for all four rooms. The daily average of time monitored by PACTS is between 1 to 5 min. It is apparent that only a very small fraction of the work done by nursing assistants were captured by PACTS. Hence, we cannot draw longitudinal trend out of the data we have collected. Nevertheless, the mean daily registered time per room was increased by more than five fold after the enabling of the lease-based mechanism.

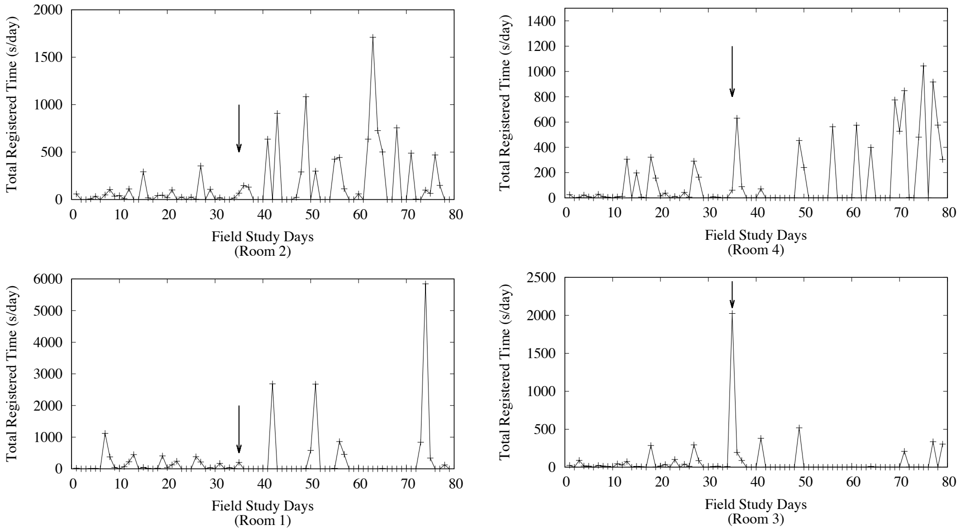

Figure 8. The total registered time per day during the entire 80-day field study in four resident rooms. The arrows in the figures indicate the day when the lease-based mechanism was activated.

Table 1. Summary of the field study result. NCA is short for non-compliant activity.

Room Number	Total of Registered Time (Minutes)	Daily Average Before/After Lease Mechanism	Total NCA Time (Minutes)	Total Number of NCAs
1	316.08 (daily mean 4.95)	3.02/18.78	16.74	427 (daily mean 5.3)
2	311.56 (daily mean 2.64)	1.56/7.11	2.04	575/7.2
3	89.16 (daily mean 1.11)	0.88/5.23	19.26	345/4.3
4	171.24 (daily mean 2.14)	1.14/6.80	30.54	620/7.8

5.2. Usability Survey Result

At the end of the field test, we asked the seven participants to complete a survey. This survey has 13 questions where the participant could choose one of four answers: Yes Much/Often; Yes Somewhat; No Not Much; and No Not At All. The participants were also encouraged to provide written comments regarding their experience. The summary of the survey result is provided in Figure 9.

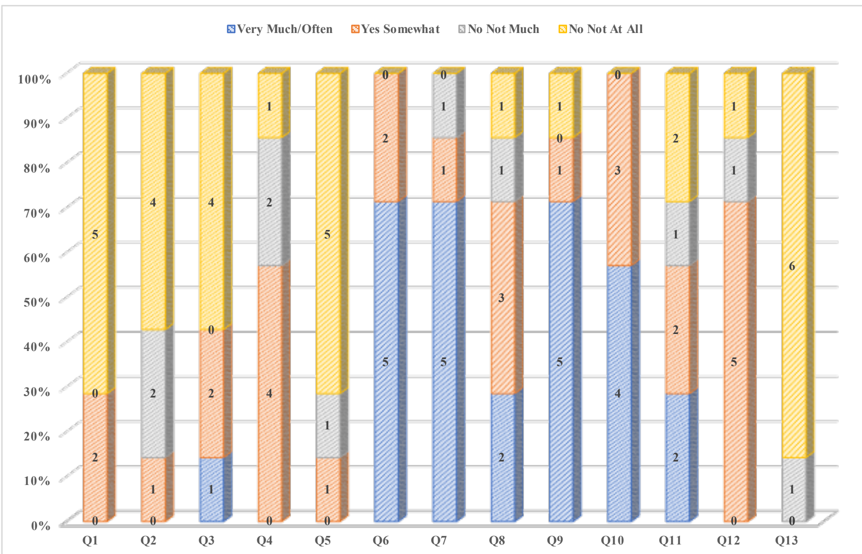

Figure 9. Survey result.

Question 1 is: "Did you mind wearing the smartwatch/carrying the smartphone?" Regarding this question, two participants indicated that to some degree they feel that it is a burden to wear the smartwatch and/or carry the smartphone, and five participants did not mind doing so. This suggests that the far majority of participants consider this system requirement acceptable. However, a non-negligible portion of participants do consider this requirement a burden. This result is inline with our observation and we have plan to address this issue.

Question 2 is: "Did wearing the smartwatch itself interfere with your job duties?" Regarding this question, only one indicated that wearing the smartwatch interfered with her job duties to some extend. The other six participants were not bothered.

Question 3 is: "Did the manual registration step interfere with your job duties?" For this question, one indicated that it did interfere very much, two reported it interfered somewhat, and four (i.e., more than half) did not find it interfere with her job duties. Again, because close to half of the participants reported that the registration step interfered with their job duties, we consider it a high priority to develop a biometric-based method to automatically register a participant with PACTS in the near future.

Question 4 is: "Did you ever have problems getting the phone to 'wake up', connect to the internet, and/or getting the watch face activated/to the correct screen?" For this question, four out of seven reported that they encountered this problem sometimes. This shows that indeed the use of a smartphone to bridge the communication between the smartwatch and the KinectServer is problematic. We have plan to switch to Android smartwatch that can connect to the KinectServer directly via WiFi to eliminate this problem.

Question 5 is: "Did the smart watch/phone maintenance (picking it up in the morning, charging, taking on and off, etc.) interfere with your job duties?" Regarding this question, only one participant indicated that these activities interfered with her job duties somewhat. Five participants felt the maintenance activities did not interfere with their job duties at all.

Question 6 is: "Did the vibrating watch help you know when you were using poor body mechanics during bedside care?" For this question, all seven participants indicated that the vibration-based alert helped them to realize that they made a mistake when they used poor body mechanics. Five of them strongly agreed that it helped. Even though it is a small sample, the survey did confirm that our design for alert delivery method is useful.

Question 7 is: "Do you think this system would be helpful when first learning good body mechanics?" For this question, only one participant indicated that PACTS is not much useful for first learning good body mechanics. Five expressed that the system is a great help in learning good body mechanics.

Question 8 is: "Do you think that using the smart watch/Kinect system resulted in any changes in your using correct body mechanics during your job tasks?" Only two participants expressed negative opinions. One indicated that "not so much" and the other one believed that it did not help at all. Five participants (62.5%) have positive opinions. Two of them believed that the system helped them very much, and three said that the system helped them somewhat.

Question 9 is: "Was the vibration a good way to be notified when you were doing something incorrectly?" This question is rather similar to question 6, but the opinion is slightly less enthusiastic with one indicating that vibration is not at all a good way to be notified. The other six participants are positive about the alert delivery method (five very positive, and one somewhat).

Question 10 is: "Did the watch ever vibrate when you thought you were moving correctly?" All seven participants reported that they have experienced false positives from our system. Four of them indicated that this occurs very often. This shows that it is challenging to design the activity detection algorithm that works in a crowded room with unpredictable orientation of the user. We plan to enhance our current algorithm with a machine-learning based approach [9,41–44] by utilizing all joints in the upper extremity.

Question 11 is: "If so, was this a problem for you?" This question is a continuation of question 10 and is relevant to all seven participants. Four participants expressed that the false positive is a problem for them. This shows that it is urgent to develop a better activity detection algorithm that minimize false positives.

Question 12 is: "Did you ever feel like you were doing incorrect movements without any feedback (vibration) from the smartwatch?" This question is about false negatives of our system. Five indicated that sometimes they experience false negatives. Again, this exposed that our system needs improvement.

Question 13 is: "Did you mind having the Kinect camera record your body positions while doing bedside care activities?" All seven participants indicated that they do not mind having the Kinect sensor record their body positions. This confirms the acceptability of PACTS by nursing assistants.

As can been from the survey result, all seven participants are very supportive in using PACTS to improve their body mechanics. The survey does confirm our observation during the field test that there are several major areas that need to be improved, particularly the manual registration and the non-compliance activity detection accuracy. The comments left by the participants are overwhelmingly positive. A few example comments are listed below:

- "It helps me to realize how I can hurt myself by moving residents wrong to transfer a resident".
- "I think it would be a good idea".
- "I think it should be repeated once it's working properly. It would be a great source for proper body mechanics".
- "It was a good experience".
- "It was nice to participate".
- "I enjoyed the opportunity but it needs a little fixing".
- "Thanks for thinking of ways to keep us safe".
- "It was different and I enjoyed it".

6. Conclusions and Future Work

In this paper, we presented the design, implementation, the result and lessons learned from a field study of PACTS. The current version of PACTS has mechanisms to support the deployment of the system into multiple rooms and can continuously track multiple nursing assistants while they go to different rooms during the same shift. Furthermore, the registration gesture recognition and non-compliant activity detection algorithm is robust again occlusions typically found in nursing homes. A main lesson that we learned from the field study is that the usability of the system must be improved significantly because PACTS is to be used by nursing assistants with no training in engineering. The introduction of the lease-based mechanism improved the usability of the system. However, the use of the smartphone as the bridge between the communication of the smartwatch and KinectServer still proved problematic. The manual registration requirement could also be improved. In future work, we plan to switch to Android smartwatch that is capable of communicating with the KinectServer using WiFi so that the smartphone will not function in the critical path of the operation of PACTS. In a lightly longer term, we aim to introduce a biometric-based auto-registration mechanism, for example, based on speaker identification, to further improve the usability of PACTS without compromising the privacy-protection of patients.

7. Patents

Part of this paper contains information included in a US patent application 20160381328, titled "Systems and Methods for Privacy-Aware Motion Tracking with Notification Feedback", where Wenbing Zhao is the inventor [45].

Acknowledgments: We thank the anonymous reviewers for their invaluable comments and suggestions on an earlier version of this paper. This study is supported in part by a grant from the Ohio Bureau of Workers' Compensation (SPC1004933944), by an award for undergraduate research, a graduate faculty travel award, a faculty research development (FRD) award, and a FRD-Internet of Things award, all from Cleveland State University. The authors received no funding to publish open-access.

Author Contributions: W.Z. conceived and designed the system; W.Z. and A.R. researched and finalized the requirement for using PACTS in nursing homes; A.R. provided instrumental suggestions on the user interface of the system; W.Z., Q.W. and N.Z. collected and analyzed the experimental results; and W.Z, Q.W. and N.Z. wrote the paper.

Conflicts of Interest: The authors declare no conflict of interest. The founding sponsors had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, and in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

PACTS	Privacy-Aware Compliance Tracking System
XML	eXtensible Markup Language
JSON	JavaScript Object Notation
HTTP	HyperText Transfer Protocol
NCA	Non-Compliant Activity
ARP	Address Resolution Protocol
DHCP	Dynamic Host Configuration Protocol
NFC	Near Field Communication

References

1. Katz, J.N. Lumbar disc disorders and low-back pain: Socioeconomic factors and consequences. *J. Bone Jt. Surg.* **2006**, *88*, 21–24.
2. Marras, W.S.; Lavender, S.A.; Leurgans, S.E.; Rajulu, S.L.; Allread, W.G.; Fathallah, F.A.; Ferguson, S.A. The Role of Dynamic Three-Dimensional Trunk Motion in Occupationally-Related Low Back Disorders: The Effects of Workplace Factors, Trunk Position, and Trunk Motion Characteristics on Risk of Injury. *Spine* **1993**, *18*, 617–628.
3. Videman, T.; Battié, M.C. Spine update: The influence of occupation on lumbar degeneration. *Spine* **1999**, *24*, 1164–1168.
4. Mohammed, S.; Singh, D.; Johnson, G.T.; Xu, P.; McCluskey, J.D.; Harbison, R.D. Evaluation of occupational risk factors for healthcare workers through analysis of the Florida Workers' Compensation Claims Database. *Occup. Dis. Environ. Med.* **2014**, *2*, 77.
5. Lun, R.; Zhao, W. A Survey of Applications and Human Motion Recognition with Microsoft Kinect. *Int. J. Pattern Recognit. Artif. Intell.* **2015**, *29*, 1555008.
6. Zhao, W.; Espy, D.D.; Reinthal, M.; Feng, H. A feasibility study of using a single Kinect sensor for rehabilitation exercises monitoring: A rule based approach. In Proceedings of the IEEE Symposium on Computational Intelligence in Healthcare and E-health, Orlando, FL, USA, 9–12 December 2014; pp. 1–8.
7. Zhao, W. A concise tutorial on human motion tracking and recognition with Microsoft Kinect. *Sci. China Inf. Sci.* **2016**, *59*, 93101.
8. Lin, J.S.; Kulic, D. Online Segmentation of Human Motion for Automated Rehabilitation Exercise Analysis. *IEEE Trans. Neural Syst. Rehabil. Eng.* **2014**, *22*, 168–180.
9. Li, Y.; Luo, X.; Wang, W.; Zhao, W. Human Action Recognition with Skeleton Data Using Extreme Learning Machine. In *Proceedings of the Chinese Intelligent Automation Conference*; Springer: Berlin, Germany, 2017; pp. 449–456.
10. Zhao, W.; Lun, R.; Espy, D.D.; Reinthal, M. Rule based realtime motion assessment for rehabilitation exercises. In Proceedings of the IEEE Symposium on Computational Intelligence in Healthcare and E-health, Orlando, FL, USA, 9–12 December 2014; pp. 133–140.
11. Zhao, W.; Lun, R.; Espy, D.D.; Ann Reinthal, M. Realtime Motion Assessment For Rehabilitation Exercises: Integration Of Kinematic Modeling With Fuzzy Inference. *J. Artif. Intell. Soft Comput. Res.* **2014**, *4*, 267–285.

12. Zhao, W.; Reinthal, M.A.; Espy, D.D.; Luo, X. Rule-Based Human Motion Tracking for Rehabilitation Exercises: Realtime Assessment, Feedback, and Guidance. *IEEE Access* **2017**, *5*, 21382–21394.
13. Hachaj, T.; Ogiela, M.R. Rule-based approach to recognizing human body poses and gestures in real time. *Multimedia Syst.* **2014**, *20*, 81–99.
14. Debar, H.; Wespi, A. Aggregation and correlation of intrusion-detection alerts. In *Recent Advances in Intrusion Detection*; Springer: Berlin, Germany, 2001; pp. 85–103.
15. Morin, B.; Mé, L.; Debar, H.; Ducassé, M. A logic-based model to support alert correlation in intrusion detection. *Inf. Fusion* **2009**, *10*, 285–299.
16. Valeur, F.; Vigna, G.; Kruegel, C.; Kemmerer, R. Comprehensive approach to intrusion detection alert correlation. *IEEE Trans. Dependable Secure Comput.* **2004**, *1*, 146–169.
17. Hravnak, M.; Edwards, L.; Clontz, A.; Valenta, C.; DeVita, M.A.; Pinsky, M.R. Defining the incidence of cardiorespiratory instability in patients in step-down units using an electronic integrated monitoring system. *Arch. Intern. Med.* **2008**, *168*, 1300–1308.
18. Chipara, O.; Lu, C.; Bailey, T.C.; Roman, G.C. Reliable clinical monitoring using wireless sensor networks: Experiences in a step-down hospital unit. In Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems, Zurich, Switzerland, 2–5 November 2010; pp. 155–168.
19. Imhoff, M.; Kuhls, S. Alarm algorithms in critical care monitoring. *Anesth. Analg.* **2006**, *102*, 1525–1537.
20. Roesch, M. Snort: Lightweight Intrusion Detection for Networks. In Proceedings of the 13th Systems Administration Conference, Seattle, WD, USA, 7–12 November 1999; Volume 99, pp. 229–238.
21. Hayes, D.L.; Levine, P.A. Pacemaker timing cycles. In *Cardiac Pacing and ICDs*; Blackwell Publishing: Malden, MA, USA, 2002; pp. 265–321.
22. Wang, P.; Kramer, A.; Estes, N.; Hayes, D.L. Timing cycles for biventricular pacing. *Pac. Clin. Electrophysiol.* **2002**, *25*, 62–75.
23. Cobelli, C.; Carson, E. *Introduction to Modeling in Physiology and Medicine*; Academic Press: Cambridge, MA, USA, 2008.
24. Sommer, R.; Paxson, V. Outside the closed world: On using machine learning for network intrusion detection. In Proceedings of the IEEE Symposium on Security and Privacy, Berkeley, CA, USA, 22–25 May 2010; pp. 305–316.
25. Becker, K.; Thull, B.; Käsmacher-Leidinger, H.; Stemmer, J.; Rau, G.; Kalff, G.; Zimmermann, H.J. Design and validation of an intelligent patient monitoring and alarm system based on a fuzzy logic process model. *Artif. Intell. Med.* **1997**, *11*, 33–53.
26. Convertino, V.A.; Moulton, S.L.; Grudic, G.Z.; Rickards, C.A.; Hinojosa-Laborde, C.; Gerhardt, R.T.; Blackbourne, L.H.; Ryan, K.L. Use of advanced machine-learning techniques for noninvasive monitoring of hemorrhage. *J. Trauma Acute Care Surg.* **2011**, *71*, S25–S32.
27. King, A.; Fortino, K.; Stevens, N.; Shah, S.; Fortino-Mullen, M.; Lee, I. Evaluation of a smart alarm for intensive care using clinical data. In Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, San Diego, CA, USA, 28 August–1 Septermber 2012; pp. 166–169.
28. Murukannaiah, P.K.; Singh, M.P. Platys: An Active Learning Framework for Place-Aware Application Development and Its Evaluation. *ACM Trans. Softw. Eng. Methodol.* **2015**, *24*, 1–33, doi:10.1145/2729976.
29. Zhao, W.; Espy, D.D.; Reinthal, M.; Ekelman, B.; Goodman, G.; Niederriter, J. Privacy-Aware Human Motion Tracking with Realtime Haptic Feedback. In Proceedings of the IEEE International Conference on Mobile Services, New York, NY, USA, 27 June–2 July 2015; pp. 446–453.
30. Zhao, W.; Lun, R.; Gordon, C.; Fofana, A.B.M.; Espy, D.D.; Reinthal, M.A.; Ekelman, B.; Goodman, G.D.; Niederriter, J.E.; Luo, X. A Human-Centered Activity Tracking System: Toward a Healthier Workplace. *IEEE Trans. Hum.-Mach. Syst.* **2017**, *47*, 343–355.
31. Zhao, W.; Wu, Q.; Padaraju, V.; Bbela, B.; Reinthal, A.; Espy, D.; Luo, X.; Qiu, T. A Privacy-Aware Compliance Tracking System for Skilled Nursing Facilities. In Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, Banff, AB, Canada, 5–9 October 2017; pp. 3568–3573.
32. Zhao, W.; Lun, R.; Gordon, C.; Fofana, A.B.M.; Espy, D.D.; Reinthal, A.; Ekelman, B.; Goodman, G.D.; Niederriter, J.E.; Luo, C.; et al. LiftingDoneRight: A Privacy-Aware Human Motion Tracking System for Healthcare Professionals. *Int. J. Handheld Comput. Res.* **2016**, *7*, 1–15.
33. Zhao, W. *Building Dependable Distributed Systems*; John Wiley & Sons: Hoboken, NJ, USA, 2014.

34. Chai, H.; Zhao, W. Towards trustworthy complex event processing. In Proceedings of the 5th IEEE International Conference on Software Engineering and Service Science, Beijing, China, 27–29 June 2014; pp. 758–761.
35. Chai, H.; Zhao, W. Byzantine Fault Tolerant Event Stream Processing for Autonomic Computing. In Proceedings of the IEEE 12th International Conference on Dependable, Autonomic and Secure Computing, Dalian, China, 24–27 August 2014; pp. 109–114.
36. Zhao, W. Optimistic Byzantine fault tolerance. *Int. J. Parallel Emerg. Distrib. Syst.* **2016**, *31*, 254–267.
37. Zhao, W. Performance optimization for state machine replication based on application semantics: A review. *J. Syst. Softw.* **2016**, *112*, 96–109.
38. Zhao, W. Fast Paxos Made Easy: Theory and Implementation. *Int. J. Distrib. Syst. Technol.* **2015**, *6*, 15–33.
39. Lun, R.; Gordon, C.; Zhao, W. Tracking the activities of daily lives: An integrated approach. In Proceedings of the Future Technologies Conference (FTC), San Francisco, CA, USA, 6–7 December 2016; pp. 466–475.
40. Lun, R.; Gordon, C.; Zhao, W. The design and implementation of a Kinect-based framework for selective human activity tracking. In Proceedings of the 2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC), Budapest, Hungary, 9–12 October 2016; pp. 002890–002895.
41. Luo, X.; Lv, Y.; Zhou, M.; Wang, W.; Zhao, W. A laguerre neural network-based ADP learning scheme with its application to tracking control in the Internet of Things. *Pers. Ubiquitous Comput.* **2016**, *20*, 361–372.
42. Luo, X.; Zhang, D.; Yang, L.T.; Liu, J.; Chang, X.; Ning, H. A kernel machine-based secure data sensing and fusion scheme in wireless sensor networks for the cyber-physical systems. *Future Gener. Comput. Syst.* **2016**, *61*, 85–96.
43. Luo, X.; Deng, J.; Wang, W.; Wang, J.H.; Zhao, W. A Quantized Kernel Learning Algorithm Using a Minimum Kernel Risk-Sensitive Loss Criterion and Bilateral Gradient Technique. *Entropy* **2017**, *19*, 365.
44. Luo, X.; Xu, Y.; Wang, W.; Yuan, M.; Ban, X.; Zhu, Y.; Zhao, W. Towards Enhancing Stacked Extreme Learning Machine With Sparse Autoencoder by Correntropy. *J. Frankl. Inst.* **2017**. doi:10.1016/j.jfranklin.2017.08.014.
45. Zhao, W. Systems and Methods for Privacy-Aware Motion Tracking with Notification Feedback. U.S. Patent App. 15/190,669, 29 December 2016.