SleeveAR: Augmented Reality for Rehabilitation using Realtime Feedback

Maurício Sousa
INESC-ID Lisboa / Técnico Lisboa / University of Lisbon
Lisboa, Portugal
antonio.sousa@ist.utl.pt

João Vieira
Técnico Lisboa / University of Lisbon
Lisboa, Portugal
felixvieira@tecnico.ulisboa.pt

Daniel Medeiros
INESC-ID Lisboa / Técnico Lisboa / University of Lisbon
Lisboa, Portugal
daniel.medeiros@tecnico.ulisboa.pt

Artur Arsénio
Universidade da Beira Interior
Covilhã, Portugal
arturarsenio@di.ubi.pt

Joaquim Jorge
INESC-ID Lisboa / Técnico Lisboa / University of Lisbon
Lisboa, Portugal
jorgej@acm.org

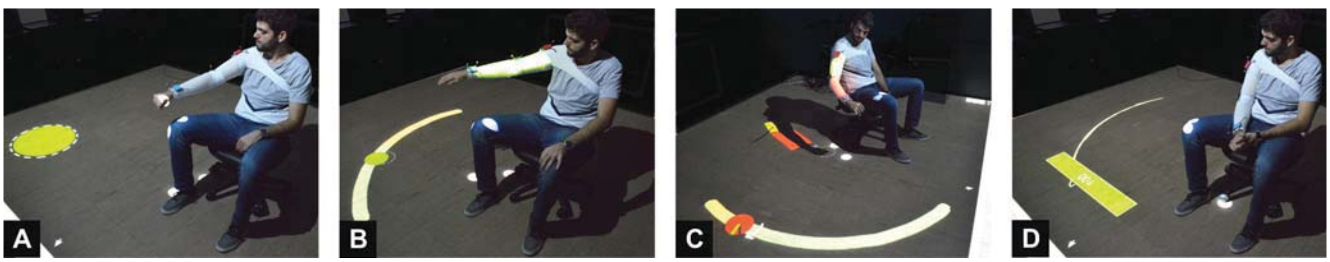

Figure 1. SleeveAR addresses new active projection-based strategies for providing user feedback during rehabilitation exercises. a) Initial position. b) Mid-performance. c) Sleeve Feedback. d) Progress report.

ABSTRACT

We present an intelligent user interface that allows people to perform rehabilitation exercises by themselves under the offline supervision of a therapist. Every year, many people suffer injuries that require rehabilitation. This entails considerable time overheads since it requires people to perform specified exercises under the direct supervision of a therapist. Therefore it is desirable that patients continue performing exercises outside the clinic (for instance at home, thus without direct supervision), to complement in-clinic physical therapy. However, to perform rehabilitation tasks accurately, patients need appropriate feedback, as otherwise provided by a physical therapist, to ensure that these unsupervised exercises are correctly executed. Different approaches address this problem, providing feedback mechanisms to aid rehabilitation. Unfortunately, test subjects frequently report having trouble to completely understand the feedback thus provided, which makes it hard to correctly execute the prescribed movements. Worse, injuries may occur due to incorrect performance of the prescribed exercises, which severely hinders recovery. SleeveAR is a novel approach to provide real-time, active feedback, using multiple projection surfaces to provide effective visualizations. Empirical evaluation shows the effectiveness of our approach as compared to traditional video-based feedback. Our experimental results show that our intelligent UI can successfully guide subjects through an exercise prescribed (and demonstrated) by a physical therapist, with performance improvements between consecutive executions, a desirable goal to successful rehabilitation.

ACM 978-1-4503-4137-0/16/03 ...$15.00.
http://dx.doi.org/10.1145/2856767.2856773

ACM Classification Keywords

H.5.2 Information Interfaces And Presentation: User Interfaces

Author Keywords

Rehabilitation; Augmented Reality; Projection-based Systems

INTRODUCTION

While supervised physical therapy is of utmost importance for the rehabilitation process, individual patient effort also plays a very significant role in their recovery. If we are to support a home rehabilitation process without the presence of a therapist, we must provide feedback to guide patients and correct them throughout their tasks. Hence, patients must be willing to learn about their condition and perform the prescribed therapeutic exercises without professional supervision. However, it is very hard to correctly perform prescribed movements without help. In the absence of proper guidance, patients may end up injuring themselves. Feedback that normally is given by a professional, by visual forms (the therapist demonstrating what to do), auditory (the therapist giving directions) or physical (the therapist applying physical force). In the absence of a therapist, it would be necessary to provide similar feedback from other sources, to support performing tasks towards desired targets.

To this end, Augmented Reality (AR) is a possible means to provide alternate feedback sources and shows a significant potential in the rehabilitation field. Additionally, tools are already available to help in developing augmented reality applications that interact with the patient's body [7]. When combined with carefully designed and customized feedback mechanisms AR can be of high value to the rehabilitation process [15]. The key idea is to provide more information to patients so that they can easily and safely execute prescribed tasks. This feedback is usually given by a therapist while undergoing physical therapy. In unsupervised settings, a different approach must be followed making sure the therapy goals are achieved and the patient correctly performs the assigned exercises. A possible approach is to take advantage of different senses by using augmented reality stimuli to help patients monitor execution of exercises. Indeed, studies have already shown that augmented reality feedback enhances the motor learning of individuals [15].

In this work, we describe SleeveAR, a novel approach that enhances patient awareness to guide them during rehabilitation exercises. SleeveAR aims at providing the means for patients to precisely replicate these exercises, especially prescribed for them by a knowledgeable health professional. Since the rehabilitation process relies on repetition of exercises during the physiotherapy sessions, our approach contributes to the correct performance of the therapeutic exercises while offering reports on the patient's progress. Furthermore, without rendering the role of the therapist obsolete, our approach builds on the notion that with proper guidance, patients can autonomously execute rehabilitation exercises. In what follows we review related work, detail the SleeveAR approach and discuss its design, informed by health professionals' feedback. We also discuss technical aspects of the implementation and present the results of its evaluation.

RELATED WORK

Our work builds on related research involving computer-assisted rehabilitation approaches and AR feedback using projection techniques. In this section we also discuss video and AR mirror approaches common in movement guidance systems.

Rehabilitation Systems

Nowadays, there are many different rehabilitation systems to help improve the recovery of patients. Many of them have different therapeutic goals and focus on specific injuries, e.g., stroke [4, 8], or limb rehabilitation [13, 6, 9]. Using these systems can greatly influence a patient's rehabilitation outside the clinic. Not only do they allow for a certain quality in performing prescribed exercises, they also enable patients to exercise in the comfortable environment of their homes. This makes it easier to stimulate and motivate them throughout the whole process [4]. Patients' effective rehabilitation is supported on three main concepts: repetition, feedback and motivation [14]. Hence, developing a rehabilitation system should be informed by these three principles and how to approach them. The repetitive nature of rehabilitation exercises can quickly become boring for a patient [12, 6, 5]. Therefore, there is a need to turn these exercises into something less tedious. When dealing with repetitive exercises, the main goal should be divided into several sub-goals. This way patients can achieve incremental success through each repetition. Furthermore, as compared to techniques where success is only achieved after completing the whole task [14], patients also report an increased motivation to improve performance. Nicolau et al. [11] used optical tracking to accurately track patient movements during rehabilitation for therapist control. Gama et al. [7] developed a rehabilitation system in which the user's position was tracked using a Microsoft Kinect. In this system, users could see themselves on the screen with overlaying targets that represented the target position. When an incorrect posture was detected (for instance, shoulders not aligned or arm not fully stretched), users were notified in real-time through visual messages. White arrows on the screen were used as visual cues to guide patients' limbs to the target. For each repetition, points were added to a score, depending on how well users performed. Klein et al.[9] focused on rehabilitating stroke victims which normally end up with one of the arms extremely debilitated. Their research focused on motivating patients to move an injured arm. Even with a small range of motion, it is important for the patient to move it to improve the recovery. The patient would see a virtual arm overlaying the injured arm, which would simulate a normal arm movement. The virtual arm position was calculated based on a few control points around the patients' shoulder and face. The results showed an enhancement of the shoulder range of motion for all test subjects. Also targeting stroke victims, Sadihov et al. [13] proposed a system to aid rehabilitation exercises via an immersive virtual environment. Tang et al. [17, 18] developed Physio@Home, a guidance system to help patients execute movements by following guidelines. The patient would see her/himself on a mirror. On top of the reflection, visual cues indicated the direction to which the arm should move. The exercises were pre-recorded by another person and then replicated by the patient.

Augmented Reality Mirrors

Mirrors allow a person to have visual feedback of his/her body. They enhance the spatial awareness which is useful to learning motor activities. AR mirrors do not necessarily require actual physical mirrors to be implemented. However, Anderson et al. [2] introduce such an approach using an actual mirror with a partially reflective layer facing the user and a diffuse layer in the back. The reflective layer provides a natural reflection while a light-projector overlays images onto the diffuse layer. The result was a mixture of the user's reflection with virtual images. Virtual mirrors can be considered an easier alternative to implement, by allowing any screen to turn into a mirror by means of a color camera. Thus we can generate virtual images on top of the reflection (for instance, for guiding purposes). Already several applications make use of augmented mirrors to guide users, be it for rehabilitation [17, 19, 9] or to support interactions outside rehabilitation [1, 3].

Projection-based Augmented Reality

Using light-projectors for AR has enabled the creation of very interesting applications. Through techniques of projection mapping, it became possible to turn any irregular surface into a projection screen. We can observe this technique applied to different objects. It is regularly used for live shows using buildings as the screen. One example is the promotion of the movie "The Tourist" where projection mapping was applied to an entire building. But it can also be applied to the human body to obtain interesting effects. Using projection mapping we can alter the perception of an object and create optical illusions. This kind of technique can bring great benefits to applications that require guiding feedback by being able to focus projection on a body part for example, just as it is necessary in rehabilitation systems. But for it to be useful, projection mapping should be interactive and applied in run-time instead of being pre-recorded like the examples above.

LightGuide [16], explored projection mapping in a innovative way. The projection was made onto users, using their body as a projection screen. Indeed, real-time visual cues were projected onto the user's hand in order to guide them through the desired movement. By projecting this information in the body part being moved, the user could keep a high level of concentration without being distracted by external factors. Different types of visual cues were developed, having in mind movements that demanded degrees of freedom over three dimensions. For each dimension a different design was planned so that the user could understand clearly in what direction should the hand move. To apply real-time projection mapping onto a moving body part, its position must be known at all time to make sure the light projector is illuminating the correct position. For this, motion tracking devices are used which enable to record the movement of, in this case, a person.

Information Feedback

Our senses are constantly at work to provide us information about our surroundings. We can think about our senses as input devices, each designed for a specific type of information. When patients are engaged on physical therapy, the therapist is constantly interacting with them. This is important in order for patients to steadily improve their performance through rehabilitation. Not only does the therapist tell them what to do but also demonstrates them movement and whenever necessary, physically corrects them. What we observe here is providing three different types of feedback to patients - audio, visual and haptic, each being interpreted by hearing, sight and touch respectively. For an automated rehabilitation system to succeed, these interactions must be simulated by other sources of feedback, in a way that the patient understands what he or she must do in the absence of the therapist. Notably, visual feedback information is often used in rehabilitation systems to communicate with users [8]. As one example of visual feedback from an AR perspective, we have the overlaying of information on an interactive mirror for the user to analyze his performance in real-time [2, 17, 19, 9, 1, 3]. Since there are multiple ways to give feedback, we can see examples where more than one are used at the same time. Combining forms of feedback can convey a better understanding of the tasks by minimizing the amount of information given visually and rely on other senses. However, if not designed with caution, a system can end up overloading the user with too much simultaneous information.

Feedback Applications

Sigrist et al. [15] suggest that different feedback types can complement each other to enhance the user comprehension. Alhamid et al. [1] introduced an interface between a user and biofeedback sensors (sensors that are able to measure physiological functions). Even though this is not aimed at rehabilitation, their approach to user interaction can be analyzed in this vein. Through this interface, users were able to access data about their bodies and health thanks to measurements conveyed by the biofeedback sensors. This system was prepared to interact with the user using multiple response interfaces, each tailored to a specific purpose. The visual interface relied on a projector that showed important messages and results from the biofeedback measurements. On the other hand, the audio interface played different kinds of music depending on the user's current state. For example, if high levels of stress were detected, relaxing music would be played to help the user relax.

One of the most common approaches to visual feedback is the augmented mirror approach already discussed. Its chief advantage is that even without overlaying virtual images, it provides users with a spatial awareness of their bodies. But we could observe other examples of augmented feedback being applied to the mirror, since a simple reflection is not enough to provide guidance. Tang et al. [17] explored two different designs for visual guidance on a mirror aimed at upper-limbs movement. Their first iteration consisted of virtual arrows that pointed at the targeted position for the user's hand. The second provided a trace of tubes placed along a path which represented the complete movement to be performed by the user's arm. However each instanced exhibited some difficulty in depth perception. This kind of visual cues has proven not to be suitable for exercises where users had to move their arms towards or away from the camera.

Anderson et al. [2] tried to provide a more detailed visual feedback by using a full virtual skeleton placed over the user reflection. In this case the goal was to mimic the skeleton's pose and hold it for a specific time. To diminish the lack of depth perception, a second tracker was placed on the user's side. Every time the system detected a large error on the z-axis, a window would appear with a side-view of both the virtual and user's skeleton. Unlike the previous approach, LightGuide [16] does not rely on interactive mirrors or screens to apply its visual feedback. By using a depth-sensor camera and a light projector, they were able to project information on the user's hand. This approach was able to guide the hand through a defined movement by projecting visual cues. All the information projected on the hand was being updated in real-time influenced by the current position given by the tracking device. The visual cues varied according to the desired direction of the movement. If the current movement only required back and forward motion, only one dimension was used. Therefore, the visual cue would only inform the user where to move his hand in the z axis through a little arrow pointing to the correct position. Two dimensional movements would combine the first visual cue by virtually painting the remaining of the hand with a color pattern. The portion of the hand closer to the desired position, would be painted with a different color than the remaining portion. They concluded that by using LightGuide, most of the users could better execute a certain movement than if they were following video instructions.

Our approach follows the work of Sodhi et al. [16] (LightGuide) and Tang et al. [17] (Physio@Home), both of them address movement guidance. But both lack performance review tools, a feature much needed during the rehabilitation process. They also assume that users always execute almost perfect movements, since the error feedback relies only in pointing to the direction of the pre-recorded exercise. Furthermore, the Physio@Home mirror metaphor provides for poor depth perception. We shall see how our interface provides for more intelligent behavior through incremental performance review and more accurate feedback.

INTERACTION DESIGN

SleeveAR deals with the specific scenario of the physical recovery of an injured arm. Our approach steps away from the mirror metaphor by augmenting, with virtual guidance information, several surfaces available to the patient, full arm and floor.

For the purpose of this research, we address basic physiotherapy exercises for upper arm and forearm, also depicted in Figure 2:

Abduction-Adduction. Horizontal movements of the arm away from or towards the center of the body.

Elevation-Depression. All arm movements above or below a vertical plane.

Flexion-Extension. Variations of the angle between the upper arm and the forearm.

Since it is required for patients to precisely perform the prescribed exercises, our approach takes extends physiotherapy sessions with exercises at home completed without direct professional supervision. However, to minimise the risk of further injuries and promote progress towards recovery, SleeveAR blends movement guidance with incremental performance reports to provide an overall awareness of the recovery process.

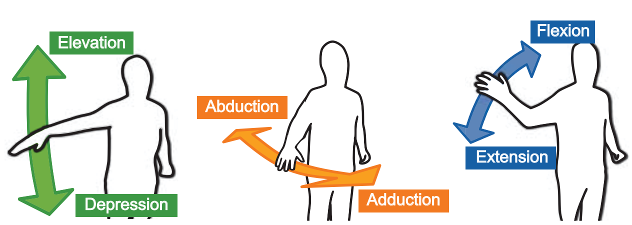

Figure 2. Physiotherapy exercises.

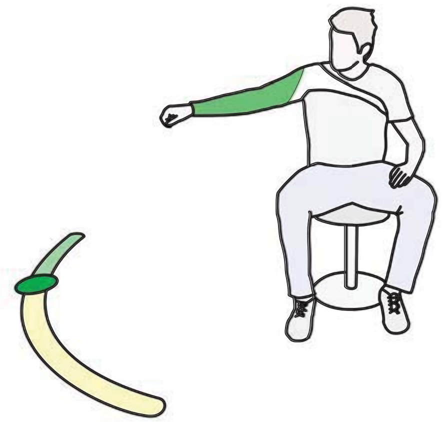

Figure 3. SleeveAR: augmented reality projections on an arm sleeve and onto the floor surface provides both movement guidance and performance feedback.

Unlike LightGuide [16], our approach takes advantage of the full arm's surface and the floor. By increasing the projection area throughout the whole arm and user's surrounding environment areas, we can successfully improve an user's awareness while a movement is being executed, as depicted in Figure 3.

In addition, the movement to be performed, is recorded by a real health professional, so that the patient can achieve a much more realistic and useful rehabilitation process. Therefore, SleeveAR main objective is to preserve the degree of similarity between the performed exercise and the one prescribed by the therapist.

SleeveAR Workflow

Our approach consists of two main concepts. First, the precise recording of the exercise being demonstrated by a personal therapist. And secondly, the ability to properly guide another person, the rehabilitation subject, during the execution of the pre-recorded exercise. Not only does this provide a vast range of possible exercises, but also it leverages on the therapist's know-how to assign adequate exercises based on a patient's condition and needs. At the same time, our method provides awareness of the rehabilitation progress to ensure the correctness of the patient's movements. In this way, a therapist can demonstrate the prescribed exercises and make sure the patient performs those correctly without requiring close supervision.

Therefore, the SleeveAR process can be divided into three main stages, as depicted in Figure 4. First, the *Recording Stage*, involves the demonstration of the exercise being recorded by the therapist. Next, the *Movement Guidance Stage*, focuses on guiding the patient to recreate the prescribed exercise as

Figure 4. SleeveAR Interaction Workflow

previously recorded. Finally, the *Performance Review Stage* provides the patient with an overview of their performance, by comparing with the original prescribed exercise.

Recording

The prescribed exercises were specifically designed for the current patient's health condition. With this in mind, we wanted to maintain this relation between a therapist and a patient, by giving therapists the power to demonstrate the prescribed exercises to the patient. Based on this demonstration, SleeveAR captures the therapist movements and stores them for a later usage. By giving the therapist the responsibility of demonstrating the exercise, we do not need to worry about the physical limitations of the patient that will use our system to recreate it. We are assuming the recorded exercise is already customized for the patient in question. Given these assumptions, SleeveAR is able to guide a patient through those exercises as best as it can.

Movement Guidance

Our approach divides the task of guiding a patient through an exercise into two stages, reaching the first *initial position* of the exercise (see Figure 1A) and *exercise performance* (Figure 1B). These two stages constitute a simple and clear process for organizing the desired activities to be performed by the patient. To successfully recreate an exercise, the user must first reach the exercise initial position, i.e., the first arm position from the recorded demonstration. To accomplish this first task, a patient must follow SleeveAR's feedback to achieve the correct arm position. After the initial position has been reached, as determined by SleeveAR, the system starts guiding the user through the recorded movement. However, it would be a very difficult task for a patient to exactly recreate the original demonstration of the exercise. To this end, SleeveAR relies on thresholds to allow some leeway in performing recorded motions. In doing so, if it were required of a patient to achieve, for example, a 90 degree arm flexion, they would not need to actually attain the exact angle. Indeed it is enough for them to achieve a flexion within the specified tolerance, which is consistent with the incremental improvement characteristic of therapy. Furthermore, SleeveAR takes into account that no two people are equal. Thus the system does not rely on absolute coordinates to guide user movements. Indeed, by normalizing the joint angles and relying on relative positions, our approach is able to accommodate people with different sized limbs, eg a 1,60m tall therapist can guide a 1,80m patient through recorded movements (or vice-versa). During the guidance task, the two different time sequences (the user's and the therapist's) are compared using the *Dynamic Time Warping* algorithm (DTW) [10], which is appropriate to map the degree of similarity between two temporal sequences that can vary both in timing and speed. Consequently, this approach overcomes the challenge of spotting movement deviations and then provide the correct feedback by dynamically discerning the differences between the two time series. Finally at the end of each exercise, SleeveAR provides an overview of the patient's performance superimposed on the original. This helps a patient in understanding what they might have done wrong and which parts of the exercise can be improved. To successfully guide a patient through their exercises while informing on their performance, we need to plan how SleeveAR interacts with its users. In the next section we describe our techniques to provide real-time and interactive feedback to the user.

Performance Review

Whenever an exercise is completed, SleeveAR provides users with a review of their performance, as shown in Figure 1D. By reviewing their exercise, a patient is able to understand how close they were to the original exercise.

Patients are informed about their performance by two different means. First, and most important, the recorded trajectory is drawn on the floor, followed by the user's most recently executed attempt. These trajectories help visualize what portions of the exercise need to be improved. Second, an overall score is calculated using the DTW algorithm and then translated to a user-friendly percentage value. With this intelligent "gamification" feature, users feel motivated to improve their score and, consequently, improve their performed exercises.

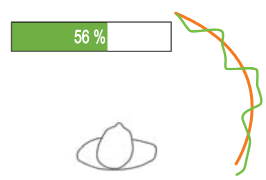

Figure 5. Example top view of the projected performance report.

In Figure 5, we depict the feedback provided after the user's movement is completed. The orange and green line represent the original trajectory and the user's attempt respectively. These are drawn on the floor. The score is depicted using a horizontal bar, which shows the degree of similarity between the performed and the recorded movements.

Real-time Feedback

Our approach provides visual feedback projected both on the user's arm and onto the floor area inside their field-of-view. It also provides audio notifications to inform the user about the end of their movement and important transitions. The motivation behind our technique is to provide users with detailed movement guidance information on the floor, while taking advantage of their peripheral vision and hearing for important notifications regarding body stance errors and cues to start/stop the exercise.

Visual Feedback

Providing useful and minimalist design was our goal when designing our visual feedback. There were some key points we wanted to address when designing it. First, the visual information had to provide the user with a representation of their *current* position, while also showing the target position. Through these representations, users must grasp easily what they need to do in order to achieve the same target position. To provide suitable feedback regarding the full arm we first applied different methods for each arm region.

Forearm. Before creating the forearm visual feedback it is important to understand what type of movement could be executed with this arm region. The forearm is connected to the upper arm by the elbow joint and its range of motion could be summarized in extension and flexing of the arm. When extending or flexing the arm, people change the elbow angle, given by the angle between the upper and forearm. To represent the

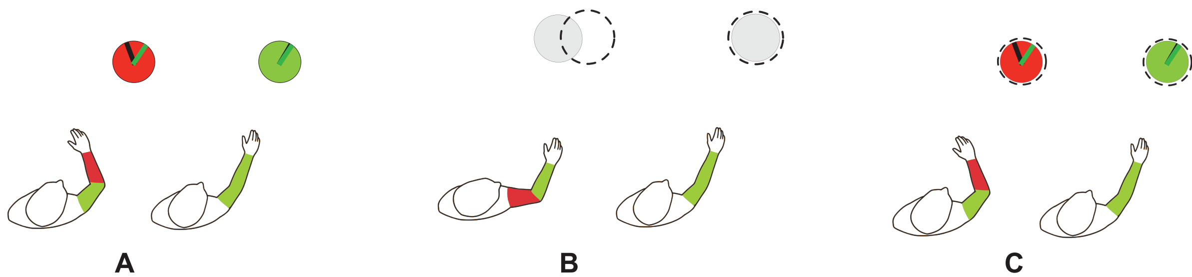

Figure 6. SleeveAR projected visual feedback. a) Forearm flexion feedback. b) Full arm positioning feedback c) Combined Full arm positioning and forearm flexion feedback.

current state, we use a black bar, as depicted in Figure 6A. Whenever the user moves their forearm, this bar moves accordingly. On the other hand, the desired forearm state is represented by the green bar. For the user to achieve this state, they must move their fore arm in order for the black bar to reach the green bar. To improve user awareness we added two additional features. Depending on the distance between the two bars, the circle color would continuously change from red (too far) to green (close enough). Also, whenever the black bar gets too far from its desired position, rotating arrows appear to warn the user that their arm is not correctly positioned.

Upper Arm. As for the upper arm, the type of movement allowed can be represented by the user's pointing direction, which is defined by the vector connecting the shoulder and the elbow joints. Once again, it is necessary to present both the current and the desired state, as shown in Figure 6B . To represent the upper arm current direction, we chose a dotted circumference. By moving the upper arm vertically or horizontally, the dotted circumference should move, away or towards the user respectively, on the floor plane.

Full Arm. Each of the two approaches above are designed to guide each arm region independently. In order to guide the user to a full arm position, we combine both, as depicted in Figure 6C. By replacing the grey circle, used to guide the upper arm, with the elbow angle circle from the forearm design, we are able to combine both simultaneously. All these techniques are used to guide the user to a specific static pose.

Additionally to the floor-projected feedback, SleeveAR projects information on the user's arm. This gives feedback to the user by representing the correctness of the current state of their arm by colors (red==incorrect / green== correct), as shown in Figure 6.

Augmented Floor. During an arm movement, we can not assume that both the upper and forearm remain in a static relation. There are situations were the arm remains fully extended throughout the movement or where the forearm relative position varies during the movement. In this case there is an elbow angle variation which means the forearm desired state is continuously changing. To this end, the feedback must then change during the movement. As for the upper arm, to help the user know where they must move it, a path is drawn showing the direction to where that section must go. The forearm changes the circle proper, while the upper arm controls the dotted circumference that must coincide with the circle. To that effect, if we move the circle throughout the movement path, we are able to continuously inform the user about the desired direction while also updating what specific elbow angle they should observe. Figure 7 illustrates an example where the user is already midway through the exercise.

Audio Feedback

Audio feedback plays an important role in timing and user notification contexts. We recognize the importance of audio to notify users about specific events. In the Recording phase, SleeveAR provides an audible notification when it actually starts recording. In this case, a countdown audio clip is used to prepare the therapist to position themselves at the initial position, before the actual recording starts. Another notification sound is played when the recording stops. As for the Movement Guidance phase, SleeveAR notifies the user whenever an exercise attempt starts. From here on, the main feedback is provided in visual form.

IMPLEMENTATION

We built a prototype in order to assess our assumptions that realtime feedback using projection-based augmented reality can yield better results than mirror-like video approaches in a realistic physical therapy scenario. The developed environment is comprised of input devices to track peoples' arm movements, while providing visual and audio cues to perform complex rehabilitation exercises, including progress report. We chose the Optitrack tracking system to implement our approach. This system relies on body markers to capture movement, but provides better accuracy than depth cameras (future higher-resolution, less noisy depth cameras may render these a viable option for use in domestic settings).

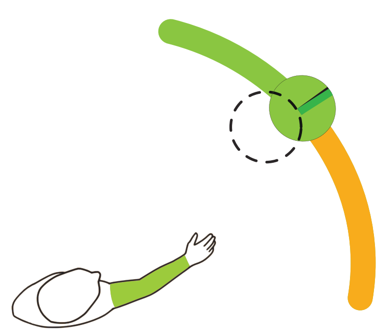

Figure 7. Movement guidance feedback

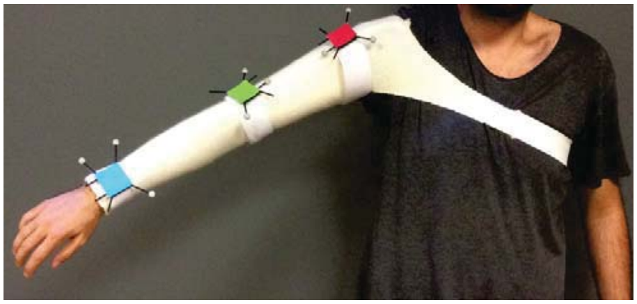

Figure 8. The Sleeve

The Sleeve

We designed a custom sleeve, as shown in Figure 8, made of wool. We employed a white colored cloth to better render light projections. To solve positioning problems, we maintain the sleeve in place using a "belt" strapped around the user's torso which greatly increased its stability. Each of the rigid bodies were still attached to a bracelet and the bracelets were stitched to the sleeve. This significantly improves the rigid bodies relative position stability, while also enabling us to tighten them more or less depending on the user's arm thickness. Another advantage of using our custom sleeve is providing a better surface to project information, due to it being white. This enable us to have a smoother and more neutral surface to project color for example. For our work, we required three different rigid bodies. Each is to be attached to a different arm location, in this case, shoulder, elbow and wrist. Having an easy way to attach and hard-to-displace method of holding our rigid bodies was vital for our work. Rigid bodies moving out of place during a movement could result in unwanted and unexpected results. Therefore, we created a better attachment method, by using a custom designed sleeve.

Given the real-time tracking information, the SleeveAR prototype then generates user feedback according to the specific exercise the user is attempting to execute. Such feedback is provided by controlling speakers to deliver audio notifications and, most importantly, by making usage of a light projector to project information both on the user's arm and floor.

Setup

Aiming for an environment favourable for the performance of rehabilitation exercises, we build our main setup in laboratory space providing the optimal area for tracking and to render projections. The SleeveAR setup is depicted in Figure 9. Multiple infrared motion cameras were installed in the ceiling for position input. Also in the ceiling, we use a commodity short-throw projector, facing downwards, to cover the maximum area surrounding the interaction space. Additionally, two audio speakers were installed as depicted in Figure 9. Our prototype handles essentially two main components:

Tracker Server. The tracker module provides position data from the user's arm captured from ten infrared motion cameras. Also, the tracker server transmits a UPD data stream of data over a local network.

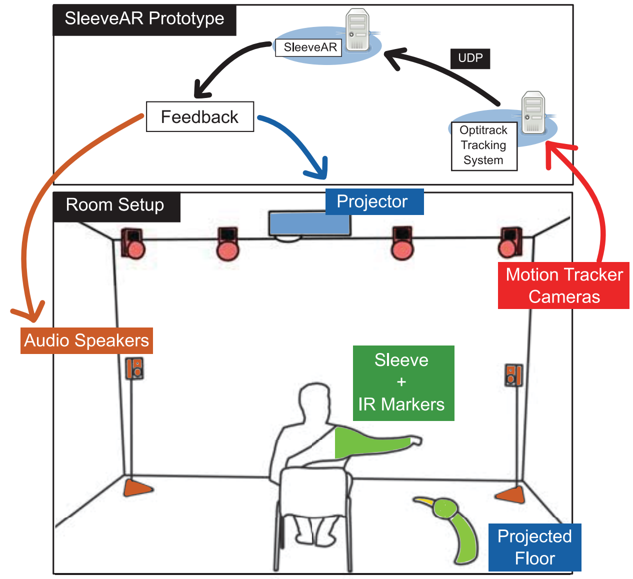

Figure 9. SleeveAR prototype's architecture and room apparatus.

SleeveAR Module. This module deals with the position input data from the tracker server, while processing it to determine the correct response, visual or audio feedback. In addition, the SeeveAR module utilizes the positional vector between the projector and the sleeve's position to determine where to project the arm's feedback techniques. Since, by calculating where the arm's shadow is cast, our prototype is able so superimpose projections on top of the sleeve surface.

EVALUATION

To evaluate SleeveAR, we intended to observe how well subjects can recreate simple arm movements just by following the feedback at their disposal. Five different exercises were defined based on such movements. Each exercise was simultaneously captured using both video and the SleeveAR's Learning component (thus the same movement is recorded both in video and in our system).

This section presents a detailed description of the user tests. We address the experimental methodology employed to test our prototype with test subjects, the category of performed tests, what metrics we used, and features of the collected sensor information. We then present the experimental results together with a critical analysis and discussion in order to achieve a better understanding about our prototype functionality and performance. Finally, as a qualitative assessment, we report the most relevant comments produced by a professional physical therapist after using our system.

Methodology

The test was divided into three stages: 1) *Introduction* 2) *Execution* and 3) *Questionnaire*. The average time spent on the test was approximately 30 minutes, of which two minutes were spent on the *Introduction*, fifteen on the *Test Execution* and the last three on the *Questionnaire*.

Exercises				
1	Abduction	Adduction		
2	Abduction	Elevation	Adduction	
3	Flexion	Abduction	Adduction	Extension
4	Elevation	Depression		
5	Abduction	Depression	Flexion	

Table 1. Arm movements and sequences during the exercises.

At the start of each session, participants received a brief explanation concerning the main goal of our prototype and were given instructions by a team member on the experiment proper. After this brief explanation, each participant performed the test. Each test session consisted of two different stages, the guidance with video and with the SleeveAR. These were shuffled at random in order to avoid experimental bias. In order to gather data for further analysis, we logged all the necessary information about the participant's movements for each exercise.

Finally, we asked participants to fill a brief questionnaire. This included questions concerning the performed tasks both with video and SleeveAR, while also providing some information about the user's profile.

The questionnaire was comprised of a list of statements scored on a 6-point Likert Scale where 1 means that users did not agree at all with a statement and 6 meant they fully agreed with it, as summarized in Tables 2 and 3.

Participants

Our subject group included 18 people, of which 14 were male and 4 were female, one of them being a physical therapist. Participants were on average 26 years old. All participants reported no physical impairment at the time of taking the tests.

Test Execution

On the task execution stage, each of the participants was asked to replicate five different rehabilitation exercises in two distinct sub-stages: *Video*, where the participant watches a video intended exercise at least two times and then, while following the video playing, the participant executes the same movement based on the video observation; and *SleeveAR*, the exactly same previously recorded exercises, now with real-time feedback. To avoid biasing the experiment, half of the participants started with *Video* while other half used *SleeveAR* first. Each exercise consisted of different movement combinations as described on Table 1.

When preparing the experiments, each exercise was simultaneously recorded with a video camera and with motion tracking devices. Under these circumstances, we made sure that the video recorded content and the data recorded by SleeveAR were the same.

In the SleeveAR phase, users would first be presented with a small tutorial which introduced each feedback component (forearm and upper arm) individually and then in combination. After the tutorial, both the SleeveAR and Video phases

It was easy to...	Video	SleeveAR
...perform the first exercise?	6 (0)	6 (0.75)
...perform the second exercise?	6 (0.75)	5.5 (1)
...perform the third exercise?	5.5 (1)	5 (2)
...perform the fourth exercise?	5.5 (1)	5 (2)
...perform the fifth exercise?	5 (1.75)	4 (1)
...follow the guidance information?	5 (1)	5 (0.75)
...see if the arm was in the right position?	5 (1.75)	5.5 (1)
...see if the arm was in the wrong position? *	6 (1.75)	6 (0.75)
...see when the exercise ended?	6 (1)	5 (1)

Table 2. User preferences: Median (Interquartile Range) for both OST and VST approaches. * indicates statistical significance. The values include Median and Inter-Quartile Range (IQR) in parentheses.

followed the same script. Each subject was allowed three attempts for each exercise, the first two served for practice and the last for measurement.

Results and Discussion

The data gathered consists of user preferences and task performance. The main objective was to address the correctness of the executed exercises. Experiments with test subjects were performed for a baseline scenario, consisting of exercise execution through video observation and a patient assisted scenario consisting of real-time feedback supported by the proposed prototype. Furthermore, this evaluation provides a formal study of our feedback techniques. Therefore, the analysis of the results is divided into a *User Preferences Overview* and *Task Performance Overview*.

User Preferences Overview

Since the values obtained from the tasks are two related-samples and come from the same population in a 6-value Likert-scale, we applied the Wilcoxon Signed Ranks test. This test was used to highlight statistically significant differences between the SleeveAR and video observation conditions.

Results from the questionnaire, as shown in Table 2, suggest that there is only one statistically relevant difference between the two tested approaches. Evidencing that, regarding user preferences, test subjects were convinced that they were capable of successfully executing all five exercises under either condition. We identified a significant statistical difference in the question - *It was easy to see if the arm was in the wrong position* - where users preferred SleeveAR to video observation (**p-value = 0.011**). This suggests that users found it easier to detect wrong movements using SleeveAR due to being constantly informed about their movement and corrected in real-time. Additionally, we observed that users seemed more interested in using SleeveAR because of its novelty and interactive feedback. Furthermore, due to the gamification provided during the performance review, the majority of users

It was easy to understand the...	Median (IQR)
...forearm feedback?	6 (0.75)
...upper arm feedback?	5.5 (1)
...full arm feedback?	5 (2)
...movement guidance feedback?	6 (1)
...arm color projection?	5 (1.5)

Table 3. Results from Widget Questionnaire

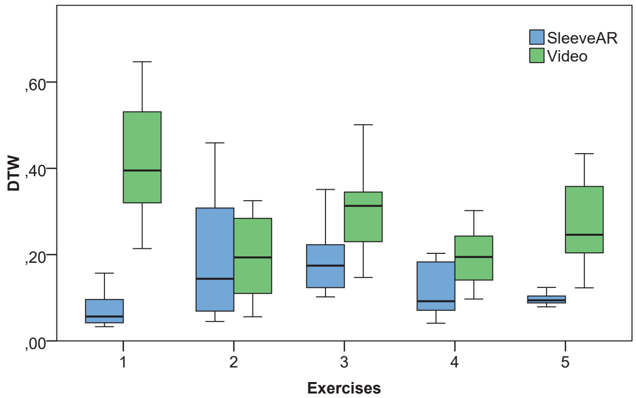

Figure 10. Average results of the Dynamic Time Warping algorithm of each exercise for both approaches (Lower values are better).

were challenging themselves to improve their score on each attempt.

Regarding visual feedback, as listed in Table 3, no statistical differences between the two approaches were observed. As for the floor feedback, some participants complained about their arms occluding visual feedback when looking down at the projections. This can be solved by positioning the floor feedback further away from the user, whenever this condition is detected. Regarding feedback on both arm and floor, some participants stated some difficulty following them simultaneously, even though these were placed in the same field of view.

Task Performance Overview

The user task performance was measured by the degree of similarity between the participants' arm paths and the original path demonstrated by the therapist. To this end we adopted again the *Dynamic Time Warping* algorithm (DTW) [10].

With the application of DTW in mind, the recorded movements can be represented as a temporal sequence of positions. One can then compare the performance values for both the proposed solution and the baseline scenario.

Due to an arm movement being divided by the upper and forearm sections, the DTW was applied to each individually, thus providing us with a more detailed set of values. This separation enables to observe if there were significant performance differences between each arm region.

The final DTW values of each exercise are the result of adding both arm regions' values. It is important to highlight that with the following results, DTW values closer to zero directly represent movements more similar to those of the original demonstration. These results show evidence that SleeveAR provided a higher similarity compared to the original exercise.

For the first exercise, one can observe in Figure 10 the test results (computed from all participants), both using the SleeveAR and just by observing the respective video. In terms of statistic values, participants achieved with SleeveAR an average DTW value of 0.114 and a Standard Deviation of 0.09, versus 0.439 and 0.165, respectively, for video observation. These results clearly suggest that SleeveAR improved participant's performance in the first exercise, i.e., they were able to re-create the original exercise better then by video observation. Based on evidence from the experimental results, similar conclusions can be drawn for the remaining four exercises.

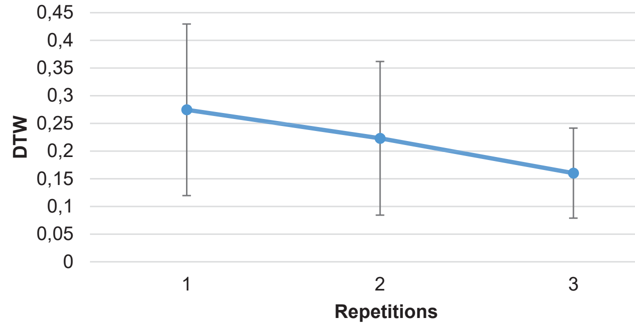

Figure 11. DTW value variation with each repetition using SleeveAR.

We applied a paired T-Student Test to data collected from each exercise on both conditions to assess the performance of our proposed solution on the user's last attempt, under a null hypothesis which stated that SleeveAR and Video observation average DTW were similar. Indeed, for all exercises the computed p-values are lower than 0.05 (0.00002, 0.00001, 0.039, 0.001 and 0.04 for the first five exercises respectively), which invalidates the null hypothesis and showing that, DTW is lower on average for SleeveAR than the video observation.

Focusing on SleeveAR results, Figure 11 presents the average DTW for each of the three trials executed by participants for each exercise. These results clearly show an improvement on a patient's performance in just a small number of repetitions. Not only the average DTW values become smaller, i.e. closer to the original, with successive repetitions, but also the standard deviation appears to diminish. Seemingly, with each repetition, participants are able to see where they performed worse in the previous attempt. These improvements on successive repetitions are statistically meaningful. Indeed, a T-student Test on the slope of the regression line, shown in Figure 11, invalidates the null hypothesis of a zero slope (i.e., performance would not increase with repetitions), with a two tailed p-value of 0.0367, lower than the significance level of 0.05.

Interview with a Physical Therapist

In addition, a professional physical therapist also tested the SleeveAR prototype, conducting the same evaluation exercises performed by the test subjects. We gathered this expert feedback in an interview as a qualitative evaluation of our technique. To recap, our main objective was to improve the guidance of subjects through pre-recorded exercises in order for them to be as close as possible to the original exercise. With this goal in mind, we wanted to investigate in which scenarios this tool might be useful in a regular physical therapy work environment. We also wanted to understand what could be missing in order to improve SleeveAR and make it a complete tool more useful in rehabilitation.

The most significant feedback is now presented, stressing both the positive and negative aspects of the proposed solution.

Missing feedback from one of the three axes. For SleeveAR feedback to be fully complete, it would need to take into account the missing axis of movement when providing real-time feedback. Since this prototype focused on guiding the arm through relatively simple movements, we did not detect this problem. But, subsequently, in the evaluation tests, we realized that it might have helped to take this into account. In the absence of checking the upper arm's rotation, SleeveAR considers some different arm poses to be the same.

Arm obstructs visibility. Occasionally, the right arm might obstruct the user's vision, making it difficult to observe the feedback being projected onto the floor. This issue could be solved by projecting the visual feedback further away from the subject, when such a situation is detected.

Increase number of tracking points in shoulder area. In physical therapy, various arm movements also focus on the shoulder area. So, it would be necessary for our sleeve to include more tracking points around the shoulder instead of only having a tracking point for the shoulder, elbow and wrist.

Potentially useful tool for patient reports. Some physical therapists follow a group of standard arm movements to initially evaluate a patient's condition. They could receive full reports with necessary data that otherwise they would have to measure physically. It could be possible to extend SleeveAR to return additional information about a patient's range of movement after executing a group of exercises. This would allow for an immediate access to information and possibly, a more precise assessment of their condition. In addition, with the possibility of recording movements and later replaying them, SleeveAR could offer a great mechanism for demonstrating to the patient, in a visual form, how much they had improved over the course of their rehabilitation, by displaying recordings of their movements.

A great tool to help a physical therapist when multitasking. While working in a physical therapy gymnasium, therapists often have to look after several patients at the same time. SleeveAR could help the therapist by reducing the number of instances they have to correct a patient and therefore, allowing them focus other patients that might need more help.

Provides a great motivation with the feedback received. The Knowledge of Performance and Results demonstrated in SleeveAR is very satisfactory and could really help in motivating patient while showing their evolution.Being able to show the patients performance by drawing their trajectories over the original exercises helps understanding which parts need improvement. Also, the real-time feedback does a great job at instantaneously showing the patient what to correct on their exercises.

CONCLUSION AND FUTURE WORK

Augmented Reality with visual feedback for rehabilitation is expected to provide a patient with improved sources of information and guidance when executing exercises outside of the controlled environment of the clinic. This is a much desirable scenario than the alternative of exercising with no feedback (and hence there is no possibility to correct a wrong execution). While the state of the art presents several solutions to provide guidance during movement execution, there is still room for improvement, and much research is needed to determine the best combination of different feedback sources. Projecting light on members to guide an individual through a movement showed promising results, but it is still difficult for patients to replicate exactly the prescribed rehabilitation exercise using this information alone.

We have described SleeveAR an intelligent user interface, which combines AR feedback and movement guidance to help rehabilitation exercises. It is intended not only to precisely guide people in performing exercises, but also to provide simple and clear awareness of the correctness or the incorrectness of the required actions, using visual and audio cues.

With SleeveAR, patients are able to formally assess feedback combinations suitable for movement guidance while solving some of the perception problems. SleeveAR also applies different feedback techniques in addition to the ones observed in the state of the art, both on the patient's body (arm and forearm) and on his surrounding floor. The ground projection shows the movements in all axes and allows the sleeve projection to continue in the patient's peripheral field of view. Furthermore, results from user tests suggest that people can replicate previously recorded movements by following multimodal feedback. Future work will address the guidance of real patients during the execution of more complex rehabilitation exercises.

We expect that full body awareness in exercises can be achieved through the usage of multiple projected surfaces (walls, furniture, and even the ceiling).This is particularly interesting to enable a larger range of exercises. Moreover, the perpendicular nature of walls can bring various one-dimensional representations of the exercises, each providing a single detailed point of view in real time. We plan as future work to determine the impact of such approaches on exercise performance.

Yet the setup and apparatus is still a bit complex, far from low cost and easy installation, to be deployed in home settings. Next efforts will focus on using commodity depth cameras to provide a more cost-effective approach. However, the current technique can be efficient, and even cost-effective, for rehabilitation gyms with multiple and concurrent therapeutic sessions.

ACKNOWLEDGMENTS

This work was partially supported by the Portuguese Foundation for Science and Technology (FCT) through the projects TECTON-3D (PTDC/EEI-SII/3154/2012), CEDAR (PTDC/EIA-EIA/116070/2009), and by national funds through FCT with reference UID/CEC/50021/2013. Daniel Medeiros would like to thank CAPES Foundation, Ministry of Education of Brazil for the scholarship grant (reference 9040/13-7). Artur Arsenio's work was partially funded by the Carnegie Mellon University (CMU) – Portuguese program through FCT, project AHA-Augmented Human Assistance, CMUP-ERI/HCI/0046/2013.

REFERENCES

1. Mohammed F. Alhamid, Mohamad Eid, and Abdulmotaleb El Saddik. 2012. A multi-modal intelligent system for biofeedback interactions. *2012 IEEE International Symposium on Medical Measurements and Applications Proceedings* (May 2012), 1–5. DOI: http://dx.doi.org/10.1109/MeMeA.2012.6226653
2. Fraser Anderson, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2013. YouMove: Enhancing Movement Training with an Augmented Reality Mirror. *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (2013), 311—-320. http://doi.acm.org/10.1145/2501988.2502045
3. T. Blum and V. Kleeberger. 2012. mirracle: An augmented reality magic mirror system for anatomy education. *Virtual Reality Short . . .* (2012), 115–116. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6180909
4. N. Alberto Borghese, Renato Mainetti, Michele Pirovano, and Pier Luca Lanzi. 2013. An intelligent game engine for the at-home rehabilitation of stroke patients. *2013 IEEE 2nd International Conference on Serious Games and Applications for Health (SeGAH)* (May 2013), 1–8. DOI:http://dx.doi.org/10.1109/SeGAH.2013.6665318
5. Grigore Burdea. 2002. Virtual Rehabilitation- Benefits and Challenges. (2002).
6. James William Burke, Michael McNeill, Darryl Charles, Philip Morrow, Jacqui Crosbie, and Suzanne McDonough. 2009. Serious Games for Upper Limb Rehabilitation Following Stroke. *2009 Conference in Games and Virtual Worlds for Serious Applications* (March 2009), 103–110. DOI: http://dx.doi.org/10.1109/VS-GAMES.2009.17
7. Alana Da Gama, Thiago Chaves, Lucas Figueiredo, and Veronica Teichrieb. 2012. Guidance and Movement Correction Based on Therapeutics Movements for Motor Rehabilitation Support Systems. *2012 14th Symposium on Virtual and Augmented Reality* (May 2012), 191–200. DOI:http://dx.doi.org/10.1109/SVR.2012.15
8. He Huang, T Ingalls, L Olson, K Ganley, Thanassis Rikakis, and Jiping He. 2005. Interactive multimodal biofeedback for task-oriented neural rehabilitation. *Engineering in Medicine and Biology Society, 2005. 27th Annual International Conference of the* (2005), 2547–2550. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1616988
9. Alexandre Klein and Gilda Aparecida De Assis. 2013. A Markeless Augmented Reality Tracking for Enhancing the User Interaction during Virtual Rehabilitation. *2013 XV Symposium on Virtual and Augmented Reality* (May 2013), 117–124. DOI: http://dx.doi.org/10.1109/SVR.2013.43
10. Joseph B Kruskal and Mark Liberman. 1983. The symmetric time-warping problem: from continuous to discrete. *Time Warps, String Edits and Macromolecules: The Theory and Practice of Sequence Comparison* (1983), 125–161.
11. Hugo Nicolau, Tiago Guerreiro, Rita Pereira, Daniel Gonçalves, and Joaquim Jorge. 2013. Computer-assisted rehabilitation: towards effective evaluation. *International Journal of Cognitive Performance Support* 1, 1 (2013), 11–26.
12. Paula Rego, PM Moreira, and LP Reis. 2010. Serious games for rehabilitation: A survey and a classification towards a taxonomy. *Information Systems and Technologies (CISTI), 2010 5th Iberian Conference on* (2010). http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5556674
13. D. Sadihov, B. Migge, and R. Gassert. 2013. Prototype of a VR upper-limb rehabilitation system enhanced with motion-based tactile feedback. *2013 World Haptics Conference (WHC)* (April 2013), 449–454. DOI: http://dx.doi.org/10.1109/WHC.2013.6548450
14. Christian Schönauer and Thomas Pintaric. 2011. Full Body Interaction for Serious Games in Motor Rehabilitation. *Proceedings of the 2Nd Augmented Human International Conference* ACM Press (2011), 1–8. http://dl.acm.org/citation.cfm?id=1959830
15. Roland Sigrist, Georg Rauter, Robert Riener, and Peter Wolf. 2013. Augmented visual, auditory, haptic, and multimodal feedback in motor learning: A review. *Psychonomic bulletin & review* 20, 1 (Feb. 2013), 21–53. DOI:http://dx.doi.org/10.3758/s13423-012-0333-8
16. Rajinder Sodhi, H Benko, and A Wilson. 2012. LightGuide: projected visualizations for hand movement guidance. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2012). http://dl.acm.org/citation.cfm?id=2207702
17. R Tang, H Alizadeh, A Tang, Scott Bateman, and Joaquim Jorge. 2014. Physio@ Home: design explorations to support movement guidance. *CHI '14 Extended Abstracts on Human Factors in Computing Systems* (2014). http://dl.acm.org/citation.cfm?id=2581197
18. Richard Tang, Anthony Tang, Xing-dong Yang Scott, and Joaquim Jorge. 2015. Physio @ Home : Exploring visual guidance and feedback techniques for physiotherapy exercises. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2015).
19. Eduardo Velloso, A Bulling, and Hans Gellersen. 2013. MotionMA: Motion modelling and analysis by demonstration. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (2013), 1309–1318. http://dl.acm.org/citation.cfm?id=2466171