Die approbierte Originalversion dieser Diplom-/Masterarbeit ist in der Hauptbibliothek der Technischen Universität Wien aufgestellt und zugänglich.
http://www.ub.tuwien.ac.at

The approved original version of this diploma or master thesis is available at the main library of the Vienna University of Technology.
http://www.ub.tuwien.ac.at/eng

FAKULTÄT FÜR !NFORMATIK
Faculty of Informatics

HomeRehab

Development of a Supportive Tool for Home Rehabilitation based on a Depth Camera

DIPLOMARBEIT

zur Erlangung des akademischen Grades

Diplom-Ingenieur

im Rahmen des Studiums

Informatik

eingereicht von

Michael Präauer
Matrikelnummer 0728095

an der
Fakultät für Informatik der Technischen Universität Wien

Betreuung: Privatdoz. Dipl.-Ing. Dr.techn Martin Kampel
Mitwirkung: Mag. Dipl.-Ing. Rainer Planinc

Wien, 11.08.2014

(Unterschrift Verfasser)

(Unterschrift Betreuung)

FAKULTÄT
FÜR !NFORMATIK
Faculty of Informatics

HomeRehab

Development of a Supportive Tool for Home Rehabilitation based on a Depth Camera

MASTER'S THESIS

submitted in partial fulfillment of the requirements for the degree of

Diplom-Ingenieur

in

Informatics

by

Michael Präauer
Registration Number 0728095

to the Faculty of Informatics
at the Vienna University of Technology

Advisor: Privatdoz. Dipl.-Ing. Dr.techn Martin Kampel
Assistance: Mag. Dipl.-Ing. Rainer Planinc

Vienna, 11.08.2014

(Signature of Author)

(Signature of Advisor)

Technische Universität Wien
A-1040 Wien ▪ Karlsplatz 13 ▪ Tel. +43-1-58801-0 ▪ www.tuwien.ac.at

Erklärung zur Verfassung der Arbeit

Michael Präauer
Stumpergasse 37/3, 1060 Wien

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit - einschließlich Tabellen, Karten und Abbildungen -, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

(Ort, Datum)	(Unterschrift Verfasser)

Acknowledgements

First of all, I would like to thank Dr. Paternostro-Sluga for her association and support in this project. Further, my thank goes to my advisor Rainer Planinc for his help and support while application development. Beside that I would like to thank my family for supporting me throughout my study and my girlfriend and friends for diverting me from the study related stress from time to time. And a special thank goes to Happy Moldan for accompanying me from study begin to end, making it a lot more enjoyable.

Abstract

Medical physical rehabilitation or physiotherapy is the most universally adopted treatment strategy for reducing human motor deficiencies caused by injuries, diseases, stroke, infarctions or old age. Physical rehabilitation is accomplished through performing precisely determined rehabilitation exercises provided by therapists, learned and performed by the patients during therapy sessions. However, due to the fact that the number of people in need of rehabilitation, as well as the medical costs are constantly increasing, the number of therapy sessions per patient is limited. This necessitates the patients to additionally perform the exercises at home, where no feedback of a therapist regarding the correctness of the exercise execution is available. The aim of this work is to evaluate the suitability of low-cost depth-sensing devices for home-based computer-assisted rehabilitation. In the course of this master's thesis, a software tool aimed on assisting the patients while performing their required rehabilitation exercises at home has been developed and a first prototype application has been implemented and evaluated. The system is based on a low-cost depth-sensing camera for tracking the patients' motions. With the information provided by this depth sensor, a basic skeletal representation of the patients' body is generated and used for calculating aberrations of the patients' motions from predetermined reference exercises provided by therapists. Through this approach, the application is able to provide the patients with the necessary feedback to perform their exercises at home correctly and further is designed to motivate the patients for accurate and frequent exercise execution. The application can be used with arbitrary rehabilitation exercises and each exercise can be adjusted according to the patients' current condition by defining specific exercise parameters. The evaluation of the application revealed promising results in regard to the performance of distinguishing between correct and incorrect exercise executions and in regard to comprehensibility, usability and motivation.

Kurzfassung

Medizinische, physische Rehabilitation oder Physiotherapie ist die am häufigsten angewandte Methode zur Linderung von körperlichen Behinderungen und Bewegungseinschränkungen die durch Unfälle, Krankheiten, Schlaganfälle, Infarkte oder hohes Alter verursacht werden können. Physische Rehabilitation wird durch die wiederholte Durchführung von exakt definierten Rehabilitationsübungen erreicht. Diese Übungen werden von Physiotherapeuten zur Verfügung gestellt und üblicherweise von den Patienten in Therapiesitzungen erlernt und dort gemeinsam mit dem Therapeuten durchgeführt. Aufgrund der Tatsache, dass sowohl die Anzahl an Personen die auf Rehabilitation angewiesen sind, als auch die medizinischen Kosten stetig steigen, ist die Anzahl an Therapiesitzungen pro Patient beschränkt. Aus diesem Grund sind die Patienten gezwungen, die nötigen Rehabilitationsübungen zusätzlich zu Hause durchzuführen, wo jedoch kein Therapeut darauf achten kann, dass die Übungen von den Patienten nicht falsch durchgeführt werden. Das Ziel dieser Arbeit ist zu überprüfen, ob kostengünstige Tiefenkameras eine geeignete Technologie für computerunterstützte Rehabilitation im Heimbereich darstellen. Im Zuge dieser Diplomarbeit wurde eine Software entwickelt und implementiert, die den Patienten dabei helfen soll ihre Rehabilitationsübungen zu Hause richtig durchzuführen. Das System basiert auf einer Tiefenkamera wodurch die Bewegungen der Patienten während der Übungsdurchführung aufgezeichnet und analysiert werden können. Aus den Daten der Tiefenkamera wird eine vereinfachte Skelettstruktur des Körpers der Patienten rekonstruiert und die Bewegungen werden in Echtzeit mit vordefinierten Referenz-Rehabilitationsübungen, die von Therapeuten zur Verfügung gestellt werden, verglichen. Basierend auf den berechneten Bewegungsabweichungen generiert die Applikation visuelles Feedback über Fehler in der Übungsdurchführung der Patienten. Das System unterstützt beliebige Rehabilitationsübungen, die durch definierbare Übungsparameter an die momentane Verfassung der jeweiligen Patienten angepasst werden können. Zusätzlich soll die Applikation auch die Motivation der Patienten für eine regelmäßige und genaue Durchführung ihrer Rehabilitationsübungen zu Hause steigern. Die Ergebnisse der Evaluierung des Systems sind sehr vielversprechend und die Software ist in der Lage zwischen richtiger und falscher Übungsdurchführung zu unterscheiden. Nutzerbefragungen über die Applikation während der Evaluierung zeigten außerdem positive Ergebnisse in Bezug auf Verständlichkeit, intuitive Interaktion und Motivation.

Contents

CHAPTER 1

Introduction

1.1 Motivation

The reasons for human motor deficiencies are manifold, including stroke [14], injuries from accidents [26], infarctions [39], diseases [34] or old age [35]. According to a study published in [42], the current leading cause of long term motor disability of adults in the western world is stroke (cerebrovascular accident). The percentage of stroke surviving patients who are suffering from motor disabilities caused by the stroke is 70% and most disabilities caused by stroke are severely restricting the patients in their activities of daily living. The kinds of motor disabilities occurring after the stroke are depending on the location where the vascular defect(s) originate [42].

Another common cause of human motor deficiencies is a spinal cord injury (SCI) [9]. Lesions on the spinal cord can result in partial or full-body paralysis, causing dramatic limitations in performing everyday activities and restrict the patients in physical independence [9]. Similar consequences can occur after traumatic brain injuries (TBI) [31] or impingement syndromes on different parts of the body, caused by accidents [16]. Furthermore, several short- or long-term motion disabilities are caused by cardiac infarctions [27]. Another prevalent reason for people suffering from motor deficiencies are a variety of diseases. A disease like Multiple Sclerosis (MS) [34] or Parkinson [18] can lead to severe physical disabilities including paralyses or limitations in the patients' motion coordination.

A natural reason for human motor deficiencies is old age. As the worldwide population continuously ages, motor disabilities caused by old age are becoming a problem which is focused on increasingly. According to estimate, people at ages over 60 worldwide will reach a number of 1.2 billions by the year 2025 and 2 billions by the year 2050 [35]. Until then, the number of people in Europe at ages over 60 will equal 40% of Europe's total population and 60% of Europe's population in working age. According to these numbers, these circumstances will deeply impact the healthcare area.

Beside all the different reasons for human motor deficiencies, the different kinds of motor deficiencies are numerous too. A typical motor disability is a limitation in the range of motion of specific body parts [21]. Another common physical limitation is that patients suffer from a reduced balance ability [31]. Also the patients' overall muscle strength and the ability of fine motor control can be reduced [21].

The most universally adopted current treatment strategy for reducing these human motor deficiencies is physical medical rehabilitation or physiotherapy [42]. The goal of physical rehabilitation is reaching a maximum level of recovery or in the best case full recovery of the disabled extremities, and further achieve a full re-integration of the patients in their social and working life [31]. Rehabilitation leads to gaining functional independence in the patients' activities of daily living [13].

Physical medical rehabilitation is accomplished through performing precisely determined rehabilitation exercises provided by therapists and tailored to the needs of the patients and to their current health condition [25]. Through a constant and continuous repetition of these specific exercises, improvements of the patients' motor function and movement capabilities are expected. These specific exercises aimed on reducing the respective physical disabilities are typically learned by the patients in rehabilitation- or therapy-sessions. In these sessions, a therapist is introducing, showing and explaining the required rehabilitation exercises to the patients and afterwards supervises and supports them while they are carrying out these exercises.

1.2 Problem Statement

Due to the increasing number of people in need of rehabilitation (e.g. through the ageing population) and due to increasing medical costs, the number of therapy sessions per patient is limited [31]. Beside that, travelling to therapy sessions is time-consuming, especially in rural areas, and difficult for elderly people [45]. This requires the patients to additionally perform the rehabilitation exercises learned in the therapy sessions at home and on their own. The drawback of this circumstance is that no feedback about the correctness of the patients' exercise execution is available, due to the fact that no therapist is supervising them. Incorrectly performed rehabilitation exercises can lead to a stagnant rehabilitation process or in the worst case can even aggravate the patients' motor disabilities [52]. Furthermore, patients often lack motivation on continuously performing the required rehabilitation exercises at home which is leading to a less frequent exercise execution and further a slowed rehabilitation process [45].

First systems for home-based computer-assisted rehabilitation, aimed on solving these problems are already available. These systems are based on elaborate motion tracking approaches for tracking the patients' motions while performing their rehabilitation exercises at home. Professional marker-based motion tracking systems or systems based on multiple camera compositions are used for this purpose [52]. However, through the emerging of low-cost depth-sensing camera technology, new possibilities for home-based computer-assisted rehabilitation arise.

1.3 Aim of the Work

The aim of this work is to verify if a system based on a single low-cost depth-sensing device, as an alternative to elaborate motion tracking systems, is suitable for home-based computer-assisted rehabilitation. In the course of this master's thesis, a software tool based on such a depth-sensor and aimed on assisting patients with motor disabilities while performing their required rehabilitation exercises at home has been developed and a first prototype application called *HomeRehab* has been implemented. This application is providing the necessary feedback the patients need to perform their exercises at home correctly. With this application, various arbitrary rehabilitation exercises can be stored and played back at the patients' home. These exercises are provided by therapists through motion-files containing the motions of a correctly performed exercise and are used by the application as a reference for correct exercise execution. The therapists can additionally tailor the provided exercises to the needs of the respective patients by defining specific exercise-parameters for each exercise to be performed. After loading such an exercise with the HomeRehab application, the contained motions are analyzed and stored while the correct exercise execution is played back and shown to the patients. Subsequently, the patients are mimicking the shown exercises while being tracked by a low-cost depth-sensing camera. The application is analyzing the data of the connected depth-sensor in real-time and is comparing the motions of the patients while performing the exercise with the reference motions provided through the motion-file of the therapist. Through this comparison, the application is able to generate distinct visual feedback about the correctness of the patients' current exercise execution. Additionally, the application is guiding the patients through the whole home-rehabilitation session by providing textual descriptions about the exercises to be performed and instructions regarding the interaction with the system. Furthermore, a score for every performed exercise is calculated and visualized in regard to the patients' exercise execution performance. Through evaluating their reached scores over time, the patients can get an overview on their rehabilitation progress. The scores are additionally serving the purpose of motivating the patients for a more accurate exercise execution and for a frequent use of the application, which further leads to a faster rehabilitation progress. The application has been developed and evaluated in association with a medical professional [1] in the area of rehabilitation.

In literature, first projects regarding home-based computer-assisted rehabilitation based on a low-cost depth-sensor have already been published (e.g. [11], [18] and [21]). These systems though, are either aimed on boosting the patients' motivation for performing their exercises at home more frequently, or they are aimed on providing the patients with feedback about the correctness of their exercise execution while performing a specific rehabilitation exercise. However, no system can yet be found which is combining both aspects and additionally is not restricted to one specific exercise, but suitable for the use of arbitrarily defined and specifically adjustable rehabilitation exercises which the HomeRehab application is capable of.

Figure 1.1 illustrates the basic structure and working principle of the HomeRehab application. The system uses pre-determined rehabilitation exercises and adjustable exercise parameters provided by therapists, as well as the data of a connected depth-sensing camera while tracking the

---

[1] *Univ. -Prof. Dr. Tatjana Paternostro-Sluga*, chair of the *Institute of Physical Medicine and Rehabilitation* at the *Sozialmedizinisches Zentrum Ost (SMZ-Ost) - Donauspital* in Vienna

patients' exercise execution at their home, as an input and generates visual feedback in real-time, informing the patients about the correctness of their performed exercises.

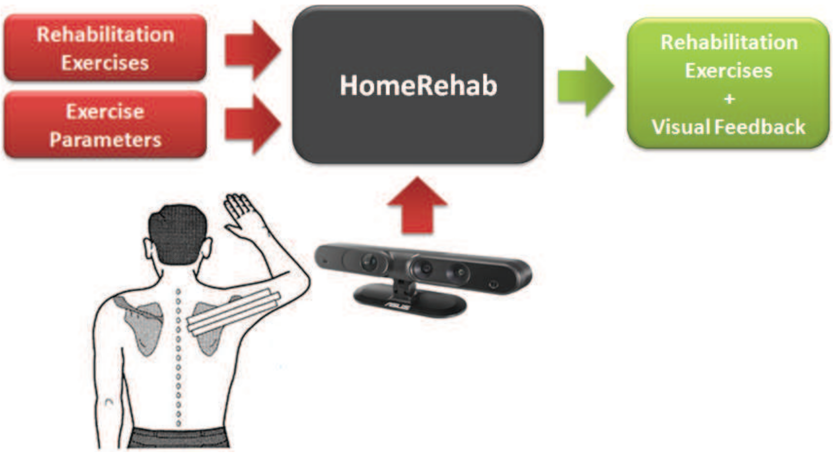

Figure 1.1: Basic structure of the HomeRehab application. Visual feedback about the correctness of the patients' rehabilitation exercise execution is generated through analyzing the data of pre-determined reference exercises and the data of a depth-sensing camera tracking the patients.

1.4 Methodological Approach

Initially, existing literature about rehabilitation in general, the state-of-the-art in human motion tracking and about methods for motion analysis and comparison has been investigated. Beside computer vision-based methods, several alternative approaches for motion analysis and comparison have been examined as well. Further, studies regarding the measurement accuracy of low-cost depth-sensing cameras have been investigated. And additionally, studies on existing home-based computer-assisted rehabilitation projects related to the HomeRehab project have been examined. After the requirement analysis for the application, a first test framework for communicating with the connected depth-sensing camera, and for acquiring and visualizing the data provided by the sensor, has been implemented. With the test framework, different approaches of motion comparison have been tested. The framework has then been extended with the functionality of loading pre-recorded motion-files and playing them back in a second sub-window simultaneously to the communication with the connected device. Subsequently, the core functionality of the application has been implemented and tested. This core functionality includes methods for creating a temporal sequence, acquiring and analyzing the skeletal representations of therapist and patient, calculating and storing angles between relevant limbs, comparing respective motions and detecting and processing motion-aberrations. Finally, the required visual guidance and feedback elements have been designed and added and the scoring mechanism has been implemented. The application has afterwards been evaluated through performance tests with participants and a usability survey. Additionally, throughout the development and evaluation stages of the project, meetings with the collaborating medical professional have been held and valuable suggestions and feedback have been gained and incorporated in the software.

1.5 Own Contribution

The suitability of low-cost depth-sensing cameras for home-based computer-assisted rehabilitation has been validated through developing and implementing a software tool for analyzing rehabilitation exercise executions through depth-information. For this purpose a motion comparison algorithm based on skeletal joint positions and angle calculations, considering specific exercise criteria, has been developed. For visualizing motion aberrations, color-coded visual feedback elements have been designed and a scoring mechanism has been developed and added. Additionally, to make the tool applicable to various different rehabilitation exercises, a unified exercise description has been developed, enabling therapists to define and adjust exercises arbitrarily.

1.6 Structure of the Work

The state-of-the-art in human motion tracking, computer-assisted rehabilitation and motion analysis is presented in Chapter 2. Additionally, the measurement accuracy of low-cost depth-sensors is examined in this chapter. The methodology for developing the HomeRehab application is introduced in Chapter 3. At first, the basic project foundations, including the depth-sensing camera and software frameworks used in this project, are explained. After that the actual application implementation, including the user interface design and the basic application structure is presented and the basic components of the application are examined. Moreover, the implemented prototype application is presented and all functionalities are explained and visualized through screenshots. In Chapter 4, the outcomes of the project evaluation are presented. The results of the evaluated rehabilitation session and the usability aspects regarding the application, as well as the limitations of the system are examined in this chapter. And finally, in Chapter 5, the results of the HomeRehab project are discussed and a conclusion is drawn.

CHAPTER 2

State of the Art

Since a core functionality of the HomeRehab project is to automatically track and analyze the motions of human body parts, this chapter highlights the state-of-the-art in human motion tracking by introducing three common contemporary motion tracking approaches. Further, the measurement accuracy of low-cost depth-sensing devices is examined. After that, existing projects on home-based computer-assisted rehabilitation are introduced and relations to the HomeRehab project are drawn. The presented projects are categorized into systems based on Serious Games and systems aimed on providing motion correction and feedback. Subsequently, existing approaches for the comparison of motions in different areas of use are examined and presented. And finally the chapter is summarized and the findings for the HomeRehab project are stated.

2.1 Human Motion Tracking

The computational tracking of human motions is required in different areas of use, including computer animation, sport sciences and medicine, and can be implemented through various ways [20]. However, this project is not focused on tracking the motions of a human body within an arbitrary scene but on detecting and tracking the motions of specific parts of a person's body while the other body parts are kept in static postures. The task of tracking motions of parts of the human body in the area of physical rehabilitation and therapy has been an active research topic since the 1980s [52]. Common examples of current human motion tracking systems are optical systems which are based on color cameras, systems based on Inertial sensors, or systems based on depth-sensing devices [20].

Optical marker-based motion capture (mocap) systems are typically used in scenarios related to computer animation, for example in the production of high-quality movies or in creating realistic computer games [20]. These systems provide the highest quality of motion data currently obtainable. However, such systems are expensive and difficult to set up.

A less cost-intensive alternative to optical motion tracking is provided by systems based on Inertial sensors. An Inertial sensor is providing its current orientation with respect to a global coordinate system. Through the lower costs and setup procedures, these sensors are available to a larger group of users and applicable in a wider range of scenarios like sports training or medical rehabilitation [20]. Inertial sensors additionally can be found in many modern devices like game-consoles or smartphones providing these devices with an additional input modality.

Depth-sensing devices are, like usual optical tracking systems, as well based on a camera tracking the scene. However, instead of capturing color images they provide images in which the intensity of each pixel represents the distance of a point in the scene to the sensor [20].

In the following subsections, the basic properties and working principles of these three sensor modalities for human motion tracking are presented.

2.1.1 Optical Sensors

Optical sensor systems are based on a set of calibrated and synchronized cameras [20]. These cameras are facing a *capture volume* in which one or more persons are performing the motions to be tracked. Every point of this capture volume can be seen by multiple cameras at any time and thus the system provides multiple views of the same person. Using triangulation algorithms, the 3D information of the tracked human body can then be deduced from these views [20]. Optical motion tracking can be implemented in two different ways by using a **marker-based** or a **marker-less** approach.

The basic idea behind using markers is that they are easily detectable in the images provided by the cameras which are recording the scene. Marker-based motion tracking is typically accomplished through providing the person to be tracked with a suit on which the markers are attached. These markers can either be passive or active. Passive markers are retro-reflective and illuminated by light sources placed closely next to each recording camera [20]. An example of an optical marker-based tracking system using passive markers is *VICON* [52]. This system provides high-quality motion tracking data with error rates around 1 mm. Figure 2.1 shows the *VICON* system while tracking the motions of an actor wearing a suit equipped with passive markers. The *VICON* motion tracking system has been used in several medical science projects as well (e.g. a gait analysis application for rehabilitation presented in [52]). An advantage of passive markers is that the person to be tracked is not required to wear wires or any other electronic equipment. A constraint of motion tracking systems based on passive markers though, is the illumination of the tracked scene. Passive marker-based systems are vulnerable to bright lighting conditions and thus are restricting the choice of the recording environment (e.g. indoor-use only) [20]. Beside that, all passive markers appear identical in the captured images and have to be identified by the motion tracking system.

Optical motion tracking systems based on active markers use LED lights as markers enabling them to emit light by themselves. Beside the advantage that no additional light sources are needed when using active markers, they can also include an encoded labelling in the emitted light, which facilitates identifying each individual marker in the recorded images [20]. An example of an optical motion tracking system based on active markers is *CODA* [52]. This system provides an accuracy of +-1.5 mm in X and Z dimension and +-2.5 mm in Y dimension when placed at a distance of 3 meters from the person to be tracked. The *CODA* system has been used in medical rehabilitation projects as well (e.g. a study on assessing muscle over-activity and spasticity on patients [52]).

Figure 2.1: The optical motion tracking system *VICON* while tracking the motions of an actor wearing a suit equipped with passive markers. Image taken from [1].

A drawback of marker-based optical motion tracking systems is that specific rotations of joints as well as overlapped body parts cannot be tracked properly. Beside that, an exact placement of the markers on the joints of the human body is required and the markers have to be attached in a way that they are not displaced or detached through the performed motions. Further, the attached markers can restrict the persons to be tracked in their movements [52].
Marker-less optical motion tracking systems estimate the position and motions of the tracked persons without the need of wearing markers through the use of computer vision in a (multiple) camera setup. Marker-less tracking of parts of the human body can either be accomplished through 2-dimensional or 3-dimensional approaches. A common method used in computer vision for the 3-dimensional reconstruction of a human body through the captured images of multiple cameras is called *Shape-From-Silhouette* [3]. With this method, the shape of an object (or person) is estimated through multiple 2D projections of the object captured at the same time in which only the outline of the object is considered. The shape provided by this method is called the *Visual Hull* of the object. For tracking the motions of the reconstructed human body, specific shape models representing an abstraction of a human body form are then matched to these 3-dimensional reconstructions. Common examples of such shape models in computer vision are stick figures or volumetric models. A stick figure is a representation of a skeletal structure comprising of segments and joint angles. Volumetric shape models representing a human body comprise of volumetric elements like elliptical cylinders, cones or volumetric blobs. A method for optical marker-less tracking of body parts through such a volumetric shape model is introduced in [8]. The proposed method provides a hierarchical 3D reconstruction from multiple camera views in real time and is based on the Shape-From-Silhouette approach. Through the

---

[1] Mocap in Education at Bradford University, `http://www.vicon.com`, Accessed: 2014-04-07

intersection of the 3D projections of the body silhouettes in the different views, a 3D voxel-based representation of the person to be tracked is reconstructed. The method uses a kinematic model comprising of volumetric blobs, representing a human body. Within the 3D space of the reconstructed person, the kinematic body model is matched and the motions of the reconstructed body parts are tracked.

A marker-less motion tracking approach based on a single camera and a 2-dimensional shape model is presented in [23]. The detection of the human body and the segmentation of body parts is accomplished through the use of multiple features including shape, contour and color. Background subtraction is used to detect motions in the tracked scene. A basic 2-dimensional shape model, consisting of two rectangles, is used to detect and track the torso of a person in the captured images. Additionally, the person's hands are tracked through a skin color model segmenting the foreground pixels into skin-color and non-skin-color regions and considering region size and relative positions.

A comparison of different marker-based and marker-less approaches presented in [20] reveals that the computational cost of retrieving the three-dimensional representation of the tracked person with marker-less optical tracking systems is higher in comparison to marker-based tracking, while the accuracy of the gained motion data is lower. Due to the fact that optical tracking systems are cost-intensive in acquisition and maintenance, and to the fact that an accurate attachment of markers in marker-based approaches is difficult, such systems are not suitable for a use in home-based scenarios [20].

2.1.2 Inertial Sensors

An Inertial sensor is providing information about accelerations and rotations, based on the physical principle of inertia. Inertial sensors typically consist of *accelerometers* for linear or angular acceleration measurement, *gyroscopes* (or rate-of-turn-sensors) for measuring angular velocities and a *magnetic field* sensor for measuring the vector of the earth's magnetic field. All these components today are put together in small boxes that can further be attached to an object or person [20]. For fusing the data provided by all the sensor components a *Kalman-filter* is commonly used. [52]. A Kalman-filter considers a series of measurements provided by the sensors, which include noise, and provides an estimate of the correct values by observing the provided measurements over time [52]. With the fused information from all three sensor types, an Inertial sensor is able to provide its orientation with respect to a global coordinate system.

Inertial sensors used for motion tracking adhere to the human body in order to collect movement information [52]. By attaching several Inertial sensors on the limbs of a human body, the rotational information of each limb is measured. From these rotations, relative positional information about the person's limb configuration is concluded. Thus, a basic skeletal representation of the tracked human body, as provided by optical sensor systems, can also be provided by Inertial sensor-based systems [20]. An example of a commonly used Inertial sensor is the *MT9* sensor shown in Figure 2.2(a). It provides an angular resolution of 0.05° (root mean square) with 1° static and 3° dynamic accuracy [52]. Figure 2.2(b) shows the *MT9* sensor while being used in a stroke rehabilitation study in combination with the optical motion tracking system *CODA*. Inertial sensor-based systems are used in different areas of medicine and rehabilitation, including systems for assessing parkinsonian rigidity [52], stroke rehabilitation [30] [5], falls

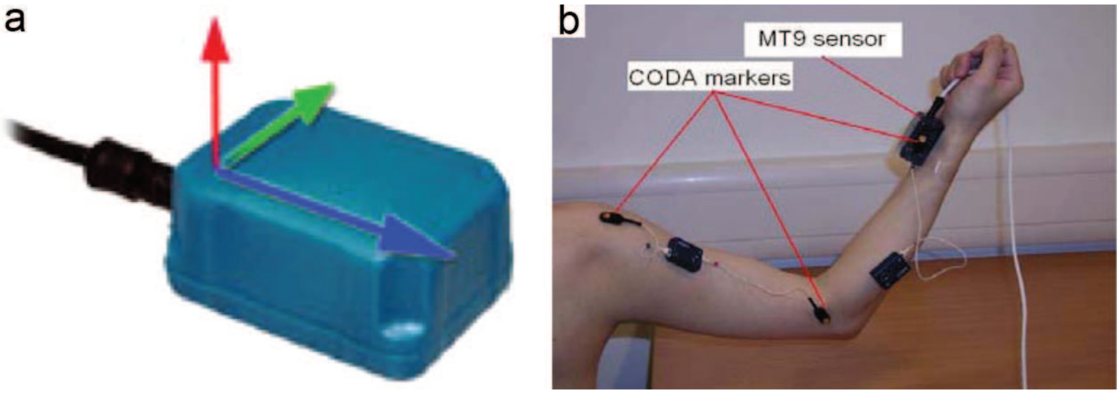

Figure 2.2: The Inertial sensor *MT9* (a) and its use in a stroke rehabilitation study (b). The proposed system is additionally including the optical motion tracking system *CODA*. Images taken from [52] (a) and [51] (b).

detection for the elderly [52] or physical activity detection on persons with chronic pulmonary disease [52].
Inertial sensors are, in contrast to optical motion tracking systems, independent from location, recording volume or illumination. Thus, they can also be operated in environments where optical tracking is not possible because of large recording volumes or uncontrollable lighting conditions. Furthermore, motion tracking with Inertial sensors is not interfered by occlusions which often can not be effectively dealt with in a home-based environment [52]. Inertial sensors nowadays are compact and lightweight, acquire high sensitivity and can be operated wirelessly. Current wireless devices have a battery live of up to 16 hours of continuous motion tracking. However, through the high number of Inertial sensors required for full body motion tracking they are less suitable for use in home application scenarios [20].

2.1.3 Depth Sensors

Another way of tracking human motions is through the use of depth-sensing cameras. Like optical tracking systems, depth-sensors are based on a camera which is capturing the scene. However, the data provided by depth-sensing cameras differs fundamentally from common color cameras. Unlike a normal RGB-camera, providing color images representing the intensities of red, green and blue light in the scene, a depth-sensing camera provides images where each pixel represents the distance of this point in the scene to the sensor [20]. From such a *depth image*, a point cloud can be deduced which is used to approximately reconstruct the 3D scene and detect and track objects within. Depth-sensing cameras can either be based on the *time-of-flight* approach or on the *structured-light* approach [20].
Time-of-flight devices determine the distance of the objects in the scene through the time the light emitted by the device takes to travel to the objects and then back to the device where a sensor is capturing the light. Since the speed of light is constant, the distance of each point in

the scene to the camera can further be calculated [20].

The second type of depth-sensors is based on structured light, where a specific point pattern is projected into the scene, distorted by the objects in the scene and then captured by the sensor [20]. Through this distortion, again, the distance of each point in the scene can be calculated.

Different computer vision approaches are used for the identification and localization of a human body in the data provided by a depth-sensor and thus for tracking the motion of specific body parts. Most common algorithms use background subtraction as a preprocessing step, where the background is assumed to be static over time to enable an extraction of the objects in the foreground [28]. One way of detecting a human body in depth information is through model-based approaches. In [50] a model-based approach for detecting and tracking persons in depth data by using a 2D head contour model and a 3D head surface model is presented. Through a distance matching algorithm scanning across the image, possible regions that may contain people are detected. Each of these regions is then examined and matched with the 3D head model. After estimating the position of the head, the whole body contour is extracted through a region growing algorithm. Another approach for detecting a human body and tracking the respective motions is through using vision-based interest point detectors on the acquired depth images. In [41], a novel interest point detector based on identifying geodesic extrema on the surface mesh is introduced. The interest points coinside with salient points of the tracked body. For detecting specific body parts, a boosted classifier then assigns local shape descriptors, extracted at the interest point locations, to body part classes.

The most popular and most widespread low-cost depth-sensing camera, based on the structured light approach, is the *Microsoft Kinect* sensor [2]. Another low-cost depth-sensing camera based on the same approach is the *Asus Xtion* [3] depth-sensor witch has been used in the HomeRehab project.

Detailed information about the depth-sensing principle of the structured light approach, as well as on the Xtion device, is provided in the section 'Project Foundations' in Chapter 3. Several studies introduce first projects based on low-cost depth-sensing cameras for medicine and rehabilitation (e.g. studies on stroke rehabilitation [47], multiple sclerosis [34] or spinal cord injuries [31]). The following section examines, if the measurement accuracy provided by such low-cost devices is sufficient for the use in home-based rehabilitation scenarios. After that, a selection of current projects regarding computer-assisted rehabilitation, based on depth-sensors as well as on other sensor modalities, is presented.

2.2 Measurement Accuracy of Low-Cost Depth Sensors

Initially, the measurement accuracy of the low-cost depth-sensing camera used in the HomeRehab project is examined to determine the precision of the provided skeletal joint positions generated by the application. In literature, several studies regarding the accuracy of the Microsoft Kinect device can be found (e.g. [9], [29], [35]). Since the Xtion PRO LIVE device

---

[2] Microsoft Kinect for Windows, `http://www.microsoft.com/en-us/kinectforwindows/`, Accessed: 2014-03-24

[3] ASUS Xtion PRO LIVE, `http://www.asus.com/Multimedia/Xtion_PRO_LIVE`, Accessed: 2014-01-30

used in this project is equipped with the exact same depth-sensor design by PrimeSense Ltd. as the Microsoft Kinect device is, the accuracy measurements in these studies apply to this depth-sensing camera as well. In [15] the motion tracking accuracy of a Kinect sensor is compared to a professional optical tracking system in a motion capture laboratory in Barcelona, comprising of 24 VICON MX3 cameras. The study examines the accuracy of the computation of joint angles by the Kinect device while tracking knee, hip and shoulder joints. The results presented show that the joint angle errors of the Kinect sensor remain beneath 10° for the knee and hip joints and between 7° and 13° for the shoulder joints while performing arbitrary motions. The higher error rates at the shoulder joints originate in the fact that these joints include three degrees of freedom which can cause occlusions while performing specific motions. The aim of this study, as well as the aim of a study comparing the Kinect to the optical tracking system *OptiTrack* [9], has been to determine if the Kinect device's motion tracking accuracy is sufficient for physical rehabilitation purposes, which has been confirmed through the acquired results in both studies. Another similar study, comparing the Kinect sensor to an optical motion tracking system, examines the accuracy of the reconstructed skeletal joint positions [38]. This paper states an average accuracy of depth reconstruction in the order of 1-4 cm at distances from the sensor of 1-4 meters. Beside the joint position accuracy, the noise of the Kinect depth-sensor while tracking a scene is dependent on the distance of the tracked object to the sensor as well as tested in [33]. At a distance of 1.2 meters, the position estimation error due to noise, averaged in x, y and z(depth) dimension is 1.3 mm, while the error at 3.5 meters distance raises to 6.9 mm [33]. This study also revealed that the positions of skeletal joints that are near to the edge of the retrieved image are constantly higher in noise [33]. Through the results of these studies it can be stated that the skeleton tracking accuracy of the low-cost depth-sensor used in the HomeRehab project is sufficient for tracking patients in home-based physical rehabilitation sessions.

2.3 Computer-Assisted Rehabilitation

The progress in computer technology, motion tracking systems and virtual reality applications offers new ways of supporting people in need of physical rehabilitation. Current studies on computer-assisted rehabilitation systems can be split into the categories *Serious Games* and *Motion Correction and Feedback*. In Serious Games aimed on physical rehabilitation, the patients are requested to move their disabled body parts in order to interact with the game and solve specific tasks in a game-like manner. These games are targeted on raising the patients motivation on constant physical activity with the disabled body parts. Through this constant activity a rehabilitation progress is expected. Systems focused on motion correction and feedback on the other hand are designed to guide the patients through specific rehabilitation exercises and provide feedback about the correctness of their exercise execution. In the following subsections, existing projects regarding these computer-assisted rehabilitation categories are presented.

2.3.1 Serious Games

In literature several studies on Serious Games focused on physical rehabilitation can be found. Since the HomeRehab project is based on a depth sensor, the focus of the presented studies lies

on systems including depth-sensing cameras. A large part of Serious Game projects based on depth-sensing cameras are working with virtual reality environments [48].
One of these projects is the application *Kinect-o-Therapy* [45] which is based on a Microsoft Kinect depth-sensor. This application includes a series of mini-games, based on conventional rehabilitation exercise routines that are targeting different parts of the patients' body. The motions of the patients are tracked by the Kinect device and passed to the virtual reality environment where the patients have to perform different tasks with their disabled limbs. These tasks range from simple actions like raising the arms as far as possible, over skill games like popping balloons for improving hand stability and hand-eye coordination, to following a given path for improving the patients' balance and coordination skills while walking. Figure 2.3 presents screenshots of four included mini-games in the *Kinect-o-Therapy* application. These exercises are called *Shoulder Exercise*, *Balloon Burst*, *Path Follower* and *Play Along*. The system is specifically designed for patients suffering from motor disabilities caused by cerebral palsy, spinal cord injury, post stroke or hereditary muscle ailments. For every game played by the patients, a score is calculated representing the patients' performance in solving the game tasks. The reached in-game scores of the patients using the system are constantly stored on a server and available to the therapist at any time via a web platform. This enables the therapist to monitor the patients' progress and provide feedback remotely. The evaluation of the project revealed

Shoulder Exercise Balloon Burst

Path Follower Play Along!

Figure 2.3: Screenshots of four included mini-games in the *Kinect-o-Therapy* application aimed on physical rehabilitation. Image taken from [45].

very promising results on the patients motivation for performing their required motions when using the application.

In [40], a similar approach with a Kinect sensor and mini-games, focused on physical rehabilitation is shown, with the difference that the gameplay of these games is automatically adapting to the patients' current performance. The patients' motions are again tracked by the Kinect device and transferred to an avatar in the virtual reality environment. One of the included games in this project is shown in Figure 2.4 where the task is to catch apples falling from a tree by navigating the avatar representing the patient from side to side. The current performance of the patients is calculated through different factors like the ratio between the number of successful trials and the total number of trials in the current game session. This information is used to constantly modify the gameplay according to the patients' limitations and current performance. Thereby the same in-game exercises can be performed by several patients with different limitations and different rehabilitation progresses and the gameplay continuously stays challenging. The project showed that through the self-adaptive approach the patients are dealing more with the application as it is more diverting than gameplays without the self-adaption.

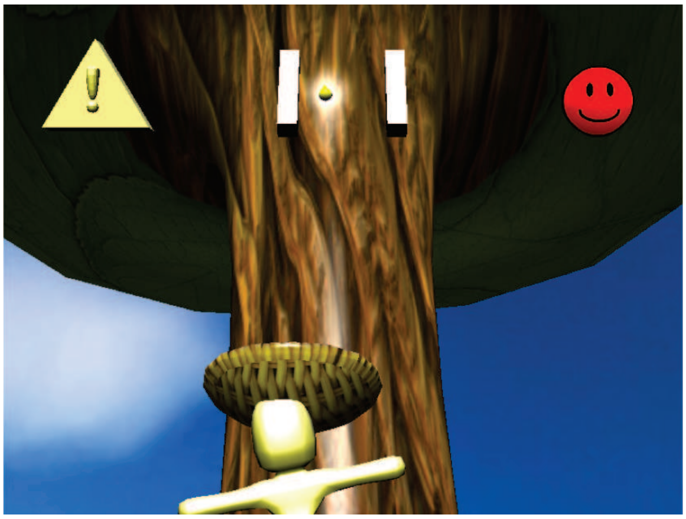

Figure 2.4: Screenshot of a self-adaptive Serious Game aimed on physical rehabilitation. The game constantly adapts the gameplay according to the current performance of the patients. Image taken from [40].

Another project including the Kinect sensor and a virtual reality environment for supporting physical rehabilitation is specifically tailored to children suffering from physical disabilities [1]. In the proposed system, the virtual environment of the open-world game *Second Life* is used as a visualization interface for the rehabilitation exercises. The idea behind the project is that disabled children, as well as their therapists interact with each other in the virtual world. The movements of the therapist while performing the required rehabilitation exercises, as well as the movements of the children while mimicking these exercises are tracked and recorded by Kinect devices and are controlling the motions of respective avatars in Second Life. Further, the chil-

dren have the possibility to load and play back the therapist-avatars, which are performing the correct exercise executions, at their home and at any time.
In the course of a project aimed on supporting physical rehabilitation of people with multiple sclerosis, an application called *REOVIEM* [34] has been developed. Again a depth-sensing camera is used to track the patients' motions while performing specific tasks tailored to multiple sclerosis. The patients are seeing themselves on a screen where additional virtual objects are overlaid. An example of a task to be performed by the patients is touching these virtual objects which are popping up on the screen with the hands or feet before the objects disappear again. The authors stated that the project has great potential on supporting patients with multiple sclerosis in their physical rehabilitation tasks.
Another Serious Game project, focused on aiding shoulder rehabilitation due to *Subacromial Impingement Syndrome* is called *The Sorcerer's Apprentice* [16]. The movements of the game avatar are again triggered by the patients' movements through a Kinect device. The game provides a freely explorable virtual environment in a fantasy setting where the patient is constantly forced to pick up items and perform specific gestures, tailored on shoulder rehabilitation, in order to proceed in the story. The results of the project evaluation regarding the patients motivation on using the application were promising and the therapists benefit from the application's possibility of recording the patients' training process for further evaluation.
A game-based rehabilitation tool focused on balance rehabilitation has been developed by members of the University of Southern California [31]. The game setting is a mine filled with gems where the patient is driving through on rails. The patient's movements are tracked by a depth-sensor and transferred to the avatar driving through the virtual mine. The task is to collect as many gems as possible by touching them with the avatar's hands while driving by. This task requires continuous stretching out of the body, the arms and the hands in all directions which is intended to improve the balance ability of the patient playing the game. A scoring system is implemented as well to motivate the patients to reach as many gems as possible for a maximum score. An initial assessment of the prototype application in association with clinicians demonstrated that the use of this prototype has great potential as a common rehabilitation tool.
A similar Kinect-based approach, but focused on patients with brain injuries is introduced in [47]. The in-game task in this application is to deflect oncoming balls with the avatar's hands (and respectively with the patients' hands) to improve cognitive and motor control skills after brain injuries. A screenshot of this application is shown in Figure 2.5. The system has been developed in collaboration with medical professionals and trialled with several patients shortly after their hospitalization. The authors state improvements in the patients response time and coordination through a constant use of the application. However, no detailed information on the time until first improvements have been encountered is given.
Beside Serious Game projects including depth-sensing cameras, systems based on haptic devices have been implemented as well. An example is a rehabilitation project for supporting patients with hand impairments which is based on a force-feedback glove [6]. This glove is transferring the motions of the patients' hands and fingers into the virtual reality environment. The patient's hand gestures are controlling a virtual hand in the virtual environment. With this virtual hand, several hand rehabilitation tasks for increasing finger force exertion and range of motion have to be performed. Additionally, specific exercises for improving the patients' hand-eye coordina-

Figure 2.5: A Serious Game aimed on improving cognitive and motor control skills after brain injuries. Image taken from [47].

tion are included as well. The results of the project evaluation again showed promising results regarding the patients' motivation in performing their specific hand-rehabilitation exercises.
Another haptic-based Serious Game aimed on wrist and arm rehabilitation is based on the *Wiimote* controller of *Nintendo's* game console *Wii* and presented in [7]. The task of the patients is to repeat specific melodies on a displayed virtual xylophone. To interact with the xylophone the patients are required to point on the respective xylophone keys with the help of the Wiimote controller. The system evaluates where the patients are pointing and is playing the respective tone. The intention of the game is that the patients are continuously performing motions with their disabled arms and wrists in order to generate the required melodies.

2.3.2 Motion Correction and Feedback

Even though there are numerous projects based on Serious Games for rehabilitation, computer-assisted rehabilitation is not constricted to game-based approaches. The projects presented in the following are less focused on games and more focused on the correction of the patients' movements while performing physical rehabilitation exercises and on providing proper feedback about the correctness of the patients' exercise execution.
A project closely related to the HomeRehab project is aimed on providing movement correction and guidance for motor rehabilitation and is introduced in [11] and [12]. A Kinect sensor is used to track and analyze shoulder abduction movements of patients performing a specific shoulder rehabilitation exercise at home. The angle between the patients' arm and trunk as well as the angle between the patients' forearm and arm are constantly calculated through the data provided by the depth-sensing camera and are further compared to the reference angles of the pre-determined correct shoulder rehabilitation exercise. The system is providing live feedback

in the form of text notifications, informing the patients whether their performed motions are correct or not. Aberrations of the patients' motions to the reference motions are monitored continuously and text notifications pop up on the screen if the aberrations are beyond a specific threshold, telling the patients where the mistake lies and what has to be changed to perform the exercise correctly. Beside that, two additional visual elements are implemented. A status bar informs the patients about the current exercise progress and a target point is displayed to inform the patients how far the arm has to be stretched to finish the exercise. Further, a score is calculated and shown, representing the patients exercise execution accuracy. Figure 2.6 shows the interface of the application with all visual feedback elements while performing one full exercise circle. The motion analysis is focused on typical mistakes occurring in this specific shoulder abduction rehabilitation exercise, for instance postural compensations. The accuracy of the angle measurement through Kinect data has been compared to the angle data of a *Goniometer* while performing the same exercise. Goniometry is a technique used by therapists for measuring the range of motion of body extremities [11]. A Goniometer has two movable arms connected by one axis which is provided with an angle measurement device, enabling the therapist to measure angles between limbs of the patients body. The project revealed that the accuracy of the Kinect depth data is sufficient for the analysis of common rehabilitation exercises. For the evaluation of the system, the shoulder abduction exercise has been performed 110 times with 50 correct exercise executions and 60 wrong ones. The system's success rate in distinguishing between

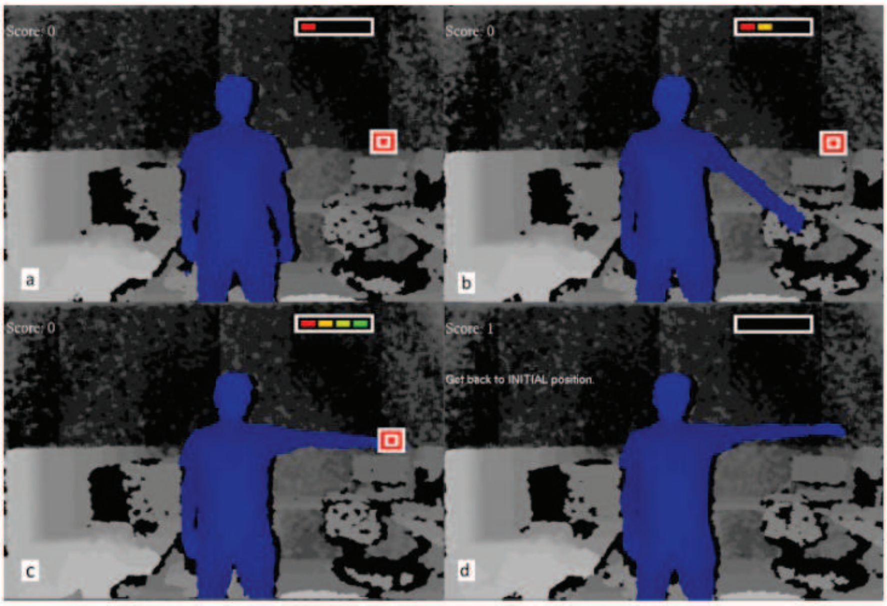

Get back to INITIAL position.

Figure 2.6: Screenshot of an application for supporting patients in performing a correct shoulder rehabilitation exercise. A full exercise circle from the initial exercise position (a) to the exercise end position (d), as well as all included visual elements (status bar, target point, scores and a text message) are shown. Image taken from [11].

correct and incorrect movements reached 100% meaning that every tested kind of aberration from the reference exercise has been detected correctly. A usability questionnaire about the application resulted positively in user satisfaction, motivation and system easiness, while letter size, information clarity and stimulus where criticized. This project, as well as the HomeRehab project, aim on avoiding wrong movements during the execution of rehabilitation exercises at home through visual feedback. However, the difference between the two projects is that this project is limited to a single, well defined rehabilitation exercise while the HomeRehab project offers a generalized software application where arbitrarily exercises can be performed.
Another Kinect-based system aiming on movement correction is focused on balance rehabilitation [18]. The application is focused on training the patients with balance disabilities in walking straight. The patients' task is to slowly walk towards the sensor on a straight line. The system analyzes the patients' gait in real-time and provides balance-related visual feedback on a screen. If the patient starts to lean to the left or right side, an arrow appears on the screen reminding the the patient to correct the stance accordingly. A screenshot of the application, including the visual feedback elements while the patient is walking towards the sensor is presented in Figure 2.7. Additionally, the patient is guided through the exercises via text messages and graphical elements. The results of the project evaluation are promising in regard to the patients' acceptance of the application and in regard of the patients' ability to use the application at home and independently.

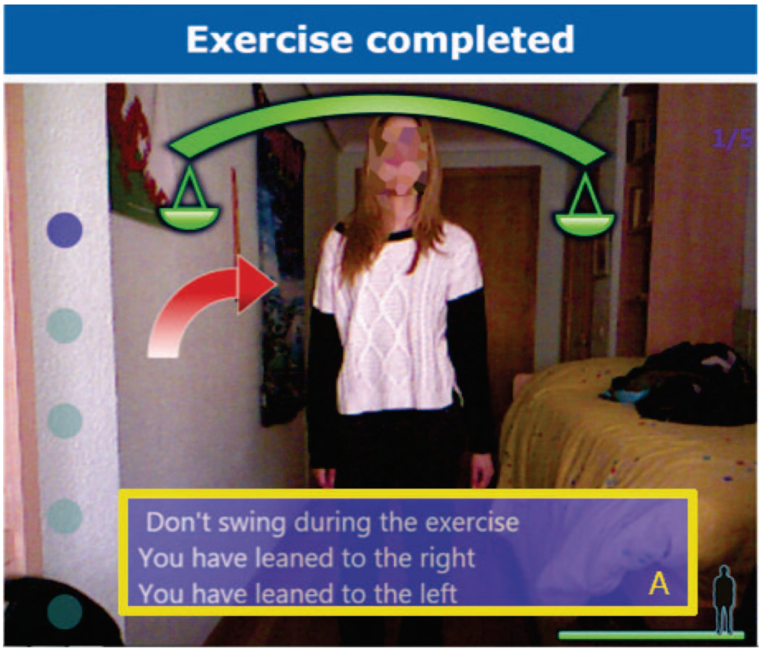

Exercise completed

Don't swing during the exercise
You have leaned to the right
You have leaned to the left

Figure 2.7: Screenshot of an application aimed on balance rehabilitation. The patient walks towards the sensor while being provided with balance-related visual feedback. Image taken from [18].

A study at the Christian University of Taiwan introduces *Kinerehab*, a depth sensor-based system aimed on motion correction for rehabilitation which is based on video and voice instructions [21]. After every video sequence the patients have to repeat the shown motions while being tracked by the depth-sensor. As soon as the patients' motions have reached a specific accuracy, the next video sequence is shown. Again, positive effects on the patients' motivation for exercise occurred through using the system.

However, computer-assisted rehabilitation projects for motion correction and feedback are not restricted to depth-sensors. Inertial sensor-based systems are common for home-based rehabilitation as well. In [37], a system design for monitoring three-dimensional movements in home rehabilitation based on Inertial sensors is presented. The prototype application uses the *Wiimote* controller of the game console *Wii* by *Nintendo* for tracking the patients' arm movements. This controller contains three one-dimensional accelerometers for covering the three dimensions X, Y and Z, as well as an infrared pointer, and can be operated wirelessly through *Bluetooth*. The system consists of two separate applications, one is patient-oriented, the other is therapist-oriented. The patient-oriented application provides the patients at home with instructions on how to perform their exercises correctly and receives the motion data from the Wiimote-controller while the patients perform the respective exercise. This data is analyzed and used to provide feedback on the screen, as well as via the LEDs located at the controller itself, whether the patient is performing the exercises correctly or not. After the exercises are completed, the motion data is stored in a XML-file together with the current date and time. With the therapist-oriented application, these XML-files can be loaded and the therapist can evaluate the performance of the patients over time, for instance to check whether the patients' range of motion has increased or not. With this information, the therapist is able to adjust the exercises to the progress of the patients and transmit the changes via the therapist-application directly to the patient-application. The project revealed the great potential of the Wiimote controller in providing therapists with the possibility of monitoring the home rehabilitation process of their patients.

A similar system has been developed in the course of a project at the University of Hannover, Germany [5]. A mobile and wireless Inertial sensor platform for motion capturing in home-based stroke rehabilitation sessions has been developed and implemented. The system is based on extensible wireless Inertial sensors featuring onboard Inertial sensor fusion, called *(IM)2SU* (Institute of Microelectronic Systems Inertial Measurement Units). These measurement units comprise tri-axial accelerometers, magnetometers and gyroscopes. The sensor fusion is implemented through a customized two step reduced state vector Kalman filter. Due to the fact that this project is focused on shoulder- and elbow-stroke-rehabilitation, the (IM)2SU sensors have to be attached to the upper and lower arm of the patients. However, the authors did not state where exactly the sensors have to be attached and if an initial calibration procedure is required. The attached sensors enable the system to compute the angles between the limbs and estimate the orientation of the patients' arm. The sensors are able to track arbitrary arm movements of the patients and communicate wirelessly with a wearable computation platform. With this computation platform the movements are analyzed and auditory feedback regarding aberrations to the correct exercise execution is generated. The basic structure of the system, including the Inertial measurement units and the wearable computation platform, is shown in Figure 2.8. The force sensor for tracking finger movements of the patients which is shown in this figure has not been part of the project and is only illustrated by the authors as a proposed future enhancement. The

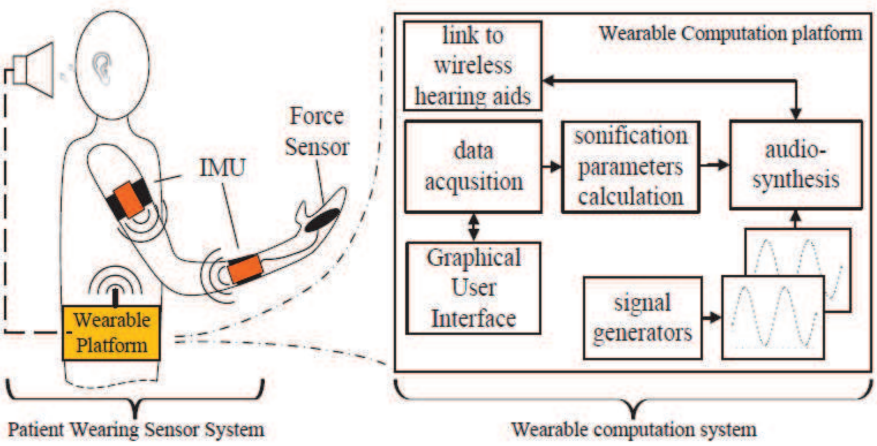

Figure 2.8: Basic structure of a stroke rehabilitation support system based on custom designed Inertial measurement units. The Inertial sensors are communicating wirelessly with a computation platform in which the patients' arm orientation is calculated and audio feedback notifying the patients about wrong exercise executions is generated. Image taken from [5].

introduced motion tracking system is extensible to up to 10 (IM)2SU sensors. The evaluation of the system revealed a reached average orientation estimation accuracy of 1.3° root mean square at a sampling rate of 50 Hz. Tests at a sampling rate of 100 Hz showed that the Inertial sensor's measurement accuracy is hardly lowered, resulting in an average orientation estimation accuracy of 1.6° root mean square.

A combined home rehabilitation approach based on both, Inertial sensors attached to the patients' body, as well as on a Kinect sensor, is presented in [4]. The Inertial sensors are used for tracking the patients' motions in performing their rehabilitation exercises while the task of the Kinect sensor in this project is to compensate the drift error of the Inertial sensors through constant re-calibration. With this approach, the high accuracy of the Inertial sensors can be maintained over time. Beside that, the data provided by the Kinect sensor is used for motion visualization.

A system based on a haptic device is called *Skyfarer* and introduced in [19]. This device comprises of an adjustable metal rig outfitted with sensors that are attached to *Thera-Bands* and free weights. The proposed system is specifically designed for supporting people in wheelchairs while performing shoulder rehabilitation exercises. Predefined shoulder exercises are shown to the patients who are subsequently asked to repeat them as accurate as possible. Through performing the shown exercises with the metal rig, the system is able to measure the patients' shoulder movements. The sensors attached to the device provide three dimensional movement data that is afterwards analyzed by the application in regard to the correctness of the exercise execution.

2.4 Methods for Comparing Motions

Beside the different possibilities of tracking human motions, there are different possibilities of comparing motions as well. One of the core functions of the HomeRehab application is to find aberrations between the tracked motions of the patients at home and the predetermined motions of the reference exercises. Thus, comparing the similarity of one motion to another. In literature, several ways of solving this problem in different areas of use can be found.

In [24] a temporal hierarchy of covariance descriptors on 3D joint locations, aimed on the correct classification of different motions is presented. Developed by members of the *Department of Computer and Systems Engineering* at the Alexandria University in Egypt, the method computes covariance matrices on the coordinates of body skeleton joints, sampled over time. The comparison of these matrices are subsequently used for human action recognition. For tracking the locations of the skeleton joints, a Kinect depth sensor is used. The proposed method has been tested on the *MSR-Action-3D-Dataset* [4] and achieved a classification rate of 90.53% meaning that over 90% of the performed motions by the participants of the study, which are representing different actions, have been classified correctly.

A motion comparison system tailored on analyzing the performance of nurses in transferring patients from a bed to a wheelchair has been developed and tested at the University of Tokyo [22]. The process of transferring a patient from a bed to a wheelchair requires a strict sequence of precisely specified actions. Two Kinect sensors are used for tracking these actions, one positioned beside the bed and tracking the scene from the side-view, the other one mounted on the ceiling, tracking the scene from above. In this project, the Kinects' depth information, as well as the RGB data is used for analyzing the movements and actions of nurse and patient. The depth information of both Kinects is used to constantly track the positions of the nurse's and patient's head, while the RGB data is used to identify and track coloured markers placed on the head, torso and ankles of both, the nurse and the patient. With all this information, the nurse's compliance with the predetermined transferring sequence can be analyzed and every step of the sequence can be evaluated by the system. For validating the system, the performance of five student nurses has been evaluated by the proposed application as well as by medical teachers, resulting in a system's evaluation accuracy of 85%.

Croitoru et al. [10] introduce a novel, non iterative 3D trajectory matching framework, invariant to translation, rotation and scale. This is achieved through a pose normalization process based on physical principles. With this normalization process, global shape signatures based on spherical coordinates can be derived and used in a matching process. These signatures are then matched through a distance measure called *LCSS* (Longest Common Sub-Sequence). The proposed method has been tested on simulated data as well as on real world data. The results have shown that the introduced method offers improved robustness and clustering accuracy compared to local measures.

Another approach for comparing motions is to use semantic descriptions which are representing human motions and actions. In [44] an ontology for modelling movement and interaction with 3D depth-sensors is presented. This ontology comprises of different features regarding user

---

[4] MSR Action Recognition Datasets and Codes, `http://research.microsoft.com/en-us/um/people/zliu/ActionRecoRsrc/`, Accessed: 2014-03-05

movement and object interaction. The goal of the presented approach is to precisely model the movements and actions a user performs through a semantic language and thus enabling automatic activity recognition or motion comparison. The results of the evaluation of the introduced ontology, composed of 164 classes, 53 object properties, 58 data properties and 93 individuals, proved its suitability for abstracting atomic gestures for an automatic activity recognition. Further, a use in the physical rehabilitation area has been stated as possible by the authors.

A system for an automatic classification of different dance gestures in real-time has been developed at the University of California [43]. The key components of the system are an angular representation of the dancer's skeleton, provided by a Kinect depth sensor, a cascaded correlation-based classifier and a distance metric to evaluate the difference in motion between an acquired gesture and a predetermined one. The system is able to correctly classify the motions of the users dancing any of the specified pre-choreographed gestures at any time while being tracked by the depth sensor. To achieve this, the system learns a statistical model that captures the nuances of a predetermined set of gesture classes and then uses the model to classify the input skeletal motion. The introduced *cascaded correlation-based maximum-likelihood multivariate classifier* introduced in this project takes into account that dancing always adheres to a musical beat which is simplifying the template matching process. The presented classifier has shown a promising robustness to various types of noise and reached an average correct classification rate of 96.9% for approximately 4-second skeletal motion recordings.

Another Kinect-based project in the area of dance gestures is presented in [2]. The aim of this project is not the classification of predetermined dance gestures but the evaluation of the overall performance of a dancer. To achieve this goal, a real-time automatic alignment algorithm has been implemented and the dancer's motions are compared through joint positions and 3D motion vectors generated through calculating joint velocities. The application is based on a virtual reality environment where the predetermined dance gestures are shown by an avatar while the user's motions are transferred to another avatar. Through this approach, the users are able to mimic shown dance choreographies and are constantly provided with scores representing their performance, calculated through the implemented motion comparison algorithms. Screenshots of the application from multiple angles and at different times are shown in Figure 2.9. The system is split into two separate applications, one for data acquisition, the other one for visualization. The paper points out that the evaluation of something as subjective as a dancer's performance is difficult. However, first tests of the system constantly resulted in higher scores for dance professionals mimicking predetermined choreographies than for amateur dancers.

A study on sports movements analysis with an Inertial sensor-based system is introduced by Jakob et.al. [26]. The task of the system is to estimate the knee angle of a person while performing knee flexion and extension movements. The wearable Inertial sensors are attached to the persons thigh and shank and aligned through a functional calibration procedure. Through estimating the relative orientations of thigh and shank, the angle of the knee joint can be calculated. In the course of the evaluation of the system, the calculated angles have been compared to the angles provided by an optical motion tracking system comprising of 8 cameras and based on passive markers which has been used as a reference. Five dynamic motions comprising of walking, jogging, running, jumps and squats have been performed by seven subjects for the evaluation. The proposed Inertial sensor-based system showed an average root mean square error of

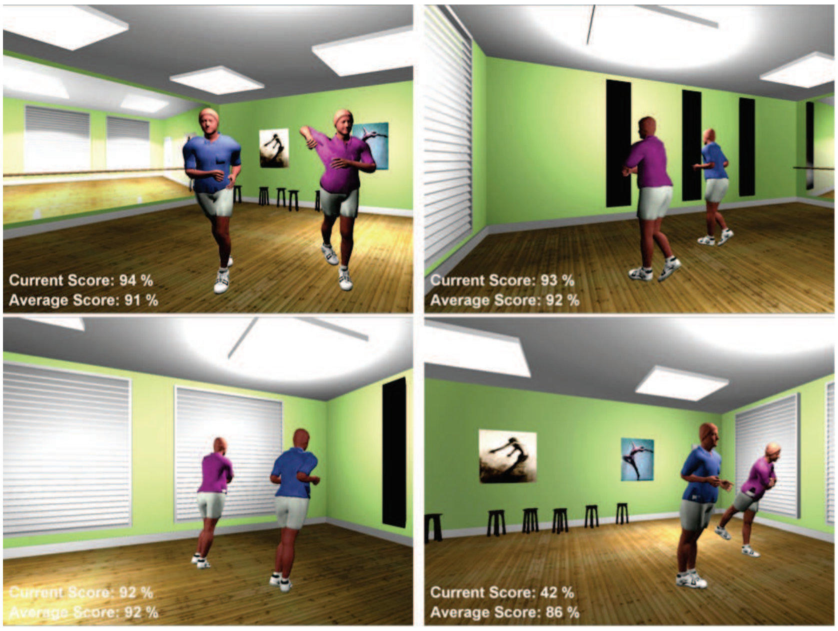

Figure 2.9: Screenshots of an application for evaluating a dancer's performance from multiple angles. On the last image an aberration of the user's dance gesture and the predetermined gesture is shown, resulting in a lower score value. Image taken from [2].

the estimated knee angle of 8.2°.
Another project focused on the measurement of angles between body limbs for motion comparison, but based on the data of a depth-sensing camera, can be found in [36]. In this study, an application called *SuperMirror* is introduced which aims on providing motion correction for ballet dancers. The system is tracking the dancers and constantly calculating the knee and hip angles formed by their limbs. For acquiring the respective angles, the limbs of the tracked skeleton are treated as 3-dimensional vectors and the angles between these vectors are calculated. The introduced system constantly compares these angles to the knee and hip angles in predetermined dance gestures and is providing visual feedback regarding aberrations of the dancers' motions from these gestures.

2.5 Summary

Through the examined state-of-the-art literature, several findings regarding the HomeRehab project can be stated. Regarding motion tracking, the suitability of a low-cost depth-sensing camera for a proper tracking of the patients' motions has been confirmed. Several studies revealed that the measurement accuracy of such devices is sufficient while the low acquisition costs of these devices enable the use in home-based scenarios. Beside that, numerous computer-assisted rehabilitation projects demonstrate their suitability for this purpose. The examined projects show that computer-assisted rehabilitation systems can increase the patients' motivation for physical exercise and that computer-assisted rehabilitation systems can vastly support patients in their exercise execution through providing proper feedback. However, the presented computer-assisted rehabilitation projects reveal that, although first approaches have been made, there is not a single application providing the possibility to define and adjust rehabilitation exercises at will which afterwards can be used for home-based rehabilitation with generated feedback. Studies regarding the comparison of motions revealed that a simple and sufficient method is through the continuous calculation and comparison of angles between the limbs of the tracked person. These angle calculations are invariant to translation, rotation and scale and cause low calculation latency. Beside that, a good collaboration of this method with the provided 3D-data of depth-sensing cameras has been shown.

CHAPTER 3

Methodology

This chapter presents the project basics and the steps taken while developing the HomeRehab prototype application, as well as an overview of the application functionalities. Initially, the project foundations are examined, including an introduction to depth-sensing cameras and to the software frameworks used in this project. Subsequently, the actual application implementation is presented. The Hard- and Software components used are listed and the design of the user interface is introduced. After that, the fundamental structure of the application and the single application components, as well as their interplay are presented and a description is given how the task of comparing motions is implemented in this project. And finally the prototype application including all program functionalities is introduced and shown through program-screenshot images.

3.1 Project Foundations

The core elements of the HomeRehab project are a depth-sensing camera as well as software frameworks for the communication with the depth sensor and for acquiring and processing the necessary data. Depth-sensing cameras in general, their working principles as well as the depth sensor used in this project are introduced in the following sections. After that, the software frameworks used for development and implementation of the HomeRehab application are examined and explained in detail.

3.1.1 Depth-Sensing Cameras

The emerging of low-cost depth-sensing camera technology started with the release of the *Microsoft Kinect* sensor [20]. Initially, the Kinect has been developed as an additional input device for Microsoft's game console *XBox-360*. Through real-time 3D human pose recognition algorithms, published by Microsoft in [46], the Kinect enables interactions between the players and the game without the need of using a controller. The interactions instead are triggered through the players' motions (*Natural Interaction*). The sensor design used in the Kinect device has been developed by the company *PrimeSense Ltd* [49]. Soon after the release of the first Kinect sensor, the computer vision society discovered that the depth-sensing technology of the Kinect could be extended far beyond gaming and at a much lower cost than traditional 3D motion tracking systems. The sensor today is used in domains like object tracking and recognition, human activity analysis, hand gesture recognition or indoor 3D mapping [28]. Beside the depth sensor itself, a Kinect device further comprises of a RGB camera and four microphones arranged in a microphone-array, as well as a motorized tilt. With these components the sensor is able to provide depth information and record RGB images and audio signals simultaneously [28]. The hardware arrangement of a Kinect device, consisting of the depth sensor including an infrared projector and an infrared camera, as well as the RGB camera and the motorized tilt is shown in Figure 3.1. The infrared projector and the infrared camera are the basic components of an approach called *Structured Light* [20] for depth-sensing, as described in the following subsection. At a distance of two meters from the Kinect sensor, it is able to resolve down to 3 mm for height and width and 1 cm for depth the in gathered depth image [49]. Other current sensors based on the Structured Light approach are the *Xtion-Pro* device by *Asus*, which has been used in developing the HomeRehab application, as well as the *Carmine* device released by *PrimeSense*. All devices incorporate the same depth sensor design developed by Primesense. The main differences of the devices lie in the device size and the image quality of the RGB camera. The Xtion and Carmine devices are more compact and lightweight than the Kinect device and provide a higher RGB image quality. Beside that, these devices are powered through the USB connection while the Kinect device requires an external power source. The Kinect device though, includes a motorized tilt for a more convenient device positioning that can be accessed and controlled remotely through the software.

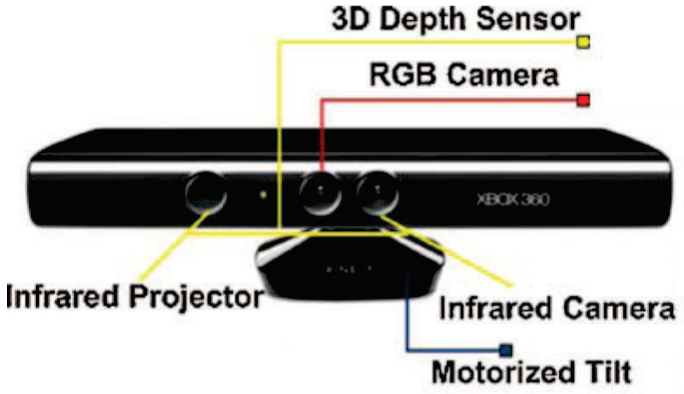

Figure 3.1: Hardware configuration of the Microsoft Kinect sensor consisting of infrared projector, RGB camera, infrared camera and motorized tilt [28].

Structured Light Approach

The central components of sensors using structured light are an infrared projector and an infrared camera. The projector and the camera are separated from each other by a fixed offset called the *baseline* [20]. The geometric relation between the projector and the camera is obtained through an offline calibration procedure applied at manufacture [28]. The actual measurement of depth can be described as a triangulation process. The infrared projector emits a single beam which is split into multiple beams through a diffraction grating. Through splitting the infrared beam, a constant pattern of speckles is projected onto the scene [29]. In Figure 3.2 such an infrared speckle pattern is shown. These speckles are invisible to the human eye as well as to RGB cameras. This pattern is then captured by the infrared camera and compared to a reference pattern which has been obtained by capturing a plane at a known distance from the sensor and stored in the memory of the device during the device's production process. When the distance of an object in the scene to the sensor is smaller or larger than the distance of this reference plane, the speckles projected on that object are shifted in the direction of the baseline. This principle is visualized in Figure 3.3. These shifts along the baseline are subsequently measured for all speckles in the captured image by a simple image correlation procedure. For each pixel of the captured image, the distance of this point in the scene to the sensor is then retrieved through these shifts [29]. To visualize the calculated distance information of the scanned scene, a *depth image* of the scene is created. Unlike a RGB image, where every pixel represents the intensities of the red, green and blue color channel, the brightness value of each pixel in a depth image represents the distance of this point in the scene to the sensor. The darker the brightness value of a pixel in this depth image, the nearer this point in the scene is to the sensor. Figure 3.4 shows

Figure 3.2: Infrared speckle pattern projected by the depth sensor. Image taken from [1].

---

[1] Kinect Raum-Scanner, Prof. Dr. Stefan Roettger, `http://schorsch.efi.fh-nuernberg.de/roettger/index.php/Projects/KinectRaum-Scanner`, Accessed: 2014-02-11

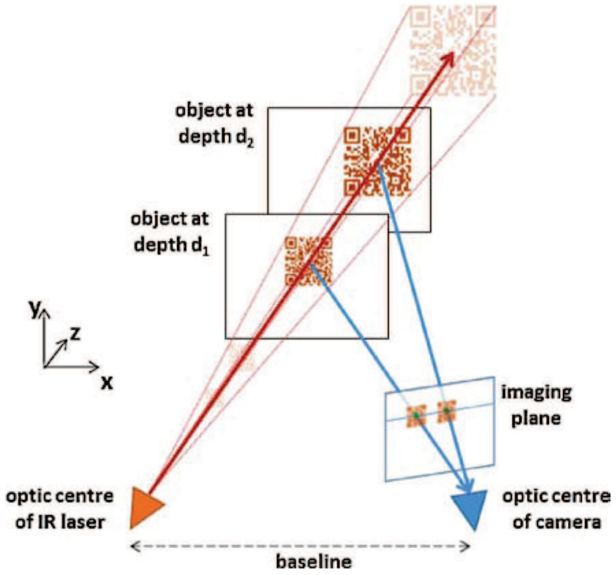

Figure 3.3: The principle of depth measurement through structured light. Image taken from [28].

a color image of an example scene (a) and its corresponding depth image (b) generated through a depth sensor scanning this scene. As can be seen by comparing the two images, areas with dark pixels in the depth image are areas near to the sensor, while light areas in the image are further away. The areas in the depth image that appear black represent tracking errors caused by occlusions, reflective surfaces or noise [20].

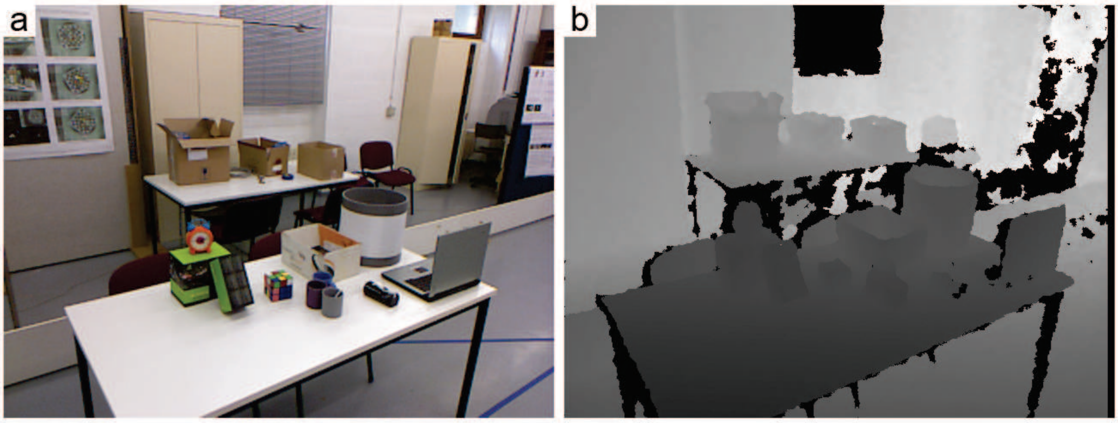

Figure 3.4: A color image of an example scene (a) and its corresponding depth image (b). Image taken from [32].

The Asus Xtion PRO LIVE

As stated before, the depth-sensing camera used in this project is the *Xtion* sensor by *Asus*. In particular the device *Xtion PRO LIVE* [2] which is equipped with a RGB camera and two microphones in addition to the depth sensor, as shown in figure 3.5. This depth-sensing camera is based on the Structured Light approach for depth measurement as well and is providing the same depth information as the Kinect device by Microsoft. The Xtion PRO LIVE is as well able to provide the depth information and the RGB images and audio streams simultaneously. The device supports several video modes for the provided depth- and RGB-streams. The RGB camera is able to provide images with a maximum resolution of 1280x1024 pixels at a framerate of 30 frames per second. The provided resolution can be reduced to a minimum of 320x240 pixels where an alternative framerate of 60 frames per second is available and can be set. The maximum provided depth image resolution is 640x480 pixels at a framerate of 30 frames per second. The depth image resolution can be reduced to 320x240 pixels as well to again be able to double the framerate to 60 frames per second. The practical operating distance of the camera for tracking human motions is within 0.8 and 3.5 meters with a field of view of 58 degrees horizontal, 45 degrees vertical and 70 degrees diagonal. The Xtion PRO LIVE depth-sensing camera is a compact device with a size of 18 cm x 3.5 cm x 5 cm and has been designed for indoor-use only. In contrast to the Xtion PRO sensor, which only supports USB 2.0, the Xtion PRO LIVE version is compatible with USB 3.0 as well. Xtion cameras have a power consumption below 2.5 W and are powered through the USB connection and therefore, in contrast to the Microsoft Kinect sensor, are not dependent on an additional power source. However, the motorized tilt of the Kinect sensor is not embedded in the ASUS Xtion devices. Both, the Kinect cameras and the Xtion cameras are being calibrated during manufacturing. Therefore the camera parameters are stored in the device's memory which enables an automatic alignment of the provided RGB and depth images [28].

Figure 3.5: The ASUS Xtion PRO LIVE depth-sensing camera comprising of infrared projector, infrared camera, RGB camera and two microphones. Image taken from [2].

---

[2] ASUS Xtion PRO LIVE, `http://www.asus.com/Multimedia/Xtion_PRO_LIVE`, Accessed: 2014-01-30

3.1.2 OpenNI Framework

Since the HomeRehab application is depending on tracking the patients motions through the depth information provided by a depth-sensing camera, a framework for accessing the respective hardware and acquiring the necessary data is needed. Microsoft is providing such a framework for free use with the official *Kinect for Windows SDK* [3]. Thus, as the name suggests, it is limited to the Microsoft Kinect sensor and additionally restricted to the Microsoft Windows operating system. As the HomeRehab project aims to be independent from the choice of the depth-sensing camera and the choice of the operating system, the Kinect for Windows SDK therefore is inappropriate. However, there are several device independent open-source frameworks available. Most commonly used are the OpenNI- and OpenKinect(LibFreeNect)-frameworks which are providing similar functionalities [28]. In this project the OpenNI [4] (Open Natural Interaction) framework is used. OpenNI has been formed by a group of companies, including PrimeSense Ltd., as a not-for-profit organization that aims to set an industry-standard framework for the interoperability of natural interaction devices. It is written and distributed under the *Apache License* and the framework's source code is freely distributed and available to the general public. The framework, provided through a software development kit, contains device drivers for all common low-cost depth-sensing cameras, libraries for accessing the data and sample projects presenting the possibilities of the framework. The latest version of the OpenNI-framework at the time of the development has been 2.2. Figure 3.6 shows the basic architecture of the OpenNI framework in version 2.2.

The OpenNI API allows any application to initialize a sensor and receive depth and RGB video

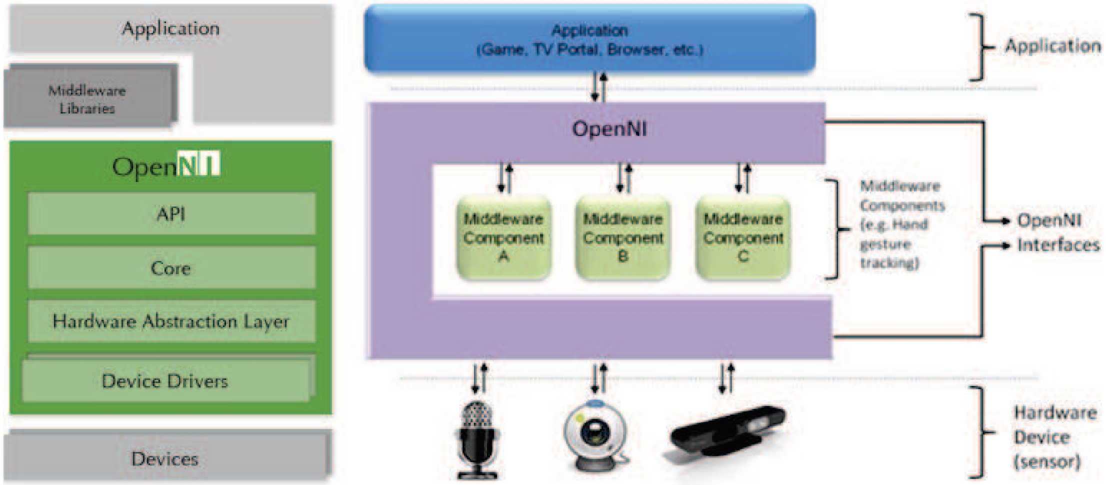

Figure 3.6: Basic architecture of the OpenNI framework in version 2.2. Images taken from [4].

---

[3] Microsoft Kinect for Windows Dev Center, `http://www.microsoft.com/en-us/kinectforwindowsdev/start.aspx`, Accessed: 2014-02-04

[4] OpenNI, `http://www.openni.org`, Accessed: 2014-02-04

streams, as well as audio streams from the microphones of the connected device. The framework includes methods for a convenient way of reading the streams frame by frame and gathering the required information. All supported depth and color video modes of the connected device can be obtained and adjusted as well. Further, several event-listener classes are included and a visualization of the RGB and depth data streams is supported. It also provides the possibility to record the gathered depth- and RGB-information from the connected sensor and transfer the data into files. These files can afterwards be loaded again with the OpenNI framework and the recorded depth and RGB information can be streamed from the files as an alternative to connecting a real device. When operated with a pre-recorded file, several playback options, like jumping to a specific frame or adjusting the playback-speed can be set as well. Beside that, several higher level middlewares for converting the raw sensor data from compliant devices to application-ready data are supported and can easily be embedded in OpenNI [49]. OpenNI is compatible with most common depth-sensing cameras and additionally provides simple multi-sensor support for operating two or more depth sensors at a time. Further, it is written in C/C++ for platform independence and the framework allows event-driven programming. The drawback of OpenNI, in contrast to the Microsoft Kinect SDK is that tracking users in sitting positions is not possible. While with the Microsoft Kinect SDK provides the possibility of tracking a user's upper body in case the lower body is not visible, the OpenNI framwork is only suitable for tracking persons in standing positions. Beside that, the Microsoft Kinect SDK supports the tracking of 20 skeletal joint positions, while the OpenNI framework is restricted to 15 as stated in the next subsection [28].

3.1.3 NiTE Middleware

As mentioned in the previous section, the OpenNI framework provides the possibility to receive raw depth data from the connected depth-sensing camera. With this data, a depth image representing the distance of each point in the scene to the sensor can be created and visualized. However, there is no further interpretation of this data. Since the goal of this project is to track human motions it is necessary to provide the application with the functionality of identifying a human body in the provided depth image. This functionality of automatically identifying human bodies in the captured depth images is provided by the *NiTE* [5] middleware. The middleware has been developed by PrimeSense Ltd. and includes specific object detection algorithms and classifiers trained on detecting a human body in the depth image and thus allowing user segmentation out of the raw depth data provided by the sensor. Other than in the previous versions of the middleware, in the version used in this project a specific calibration pose performed by the person to be tracked is not required anymore as the middleware is able to locate and track a person within the range of the sensor automatically. Further, the middleware is able to locate and segment multiple users at the same time. Other than the OpenNI framework, the source code of the NiTE middleware is not made available as open source. Figure 3.7 shows a depth image with an automatically identified and segmented human body through the NiTE middleware. Beside locating and segmenting a human body, the NiTE middleware additionally provides the functionality of deducing the basic skeletal joints of the segmented body in the depth image. The middleware

---

[5] Primesense NiTE Middleware, `http://www.primesense.com`, Accessed: 2014-02-11

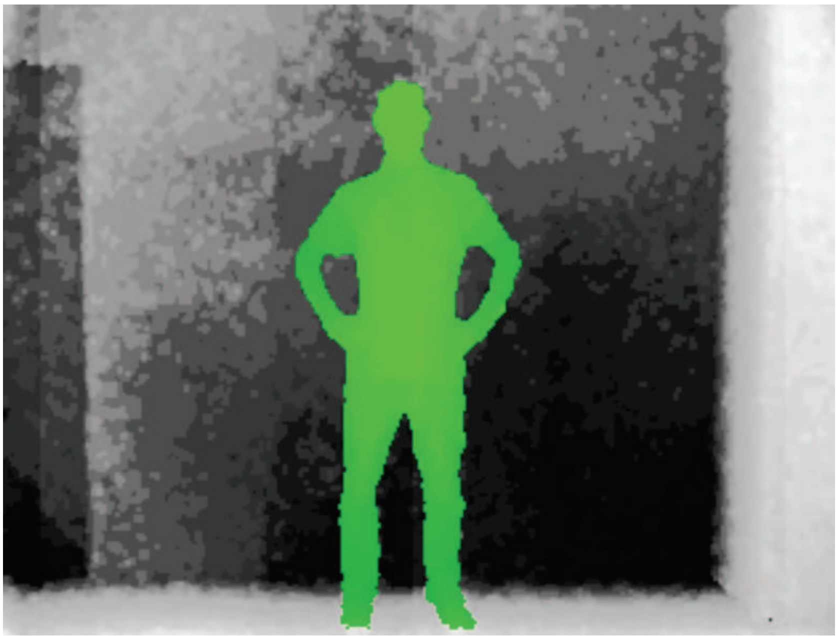

Figure 3.7: Automatic user segmentation in a captured depth image, provided by the NiTE middleware.

provides the positions and orientations of 15 skeletal joints within the segmented user, consisting of the hands, elbows and shoulders, the hips, knees and ankles, the head and the neck, as well as the center of the body's torso. By connecting these 15 joint positions a basic skeletal representation of the tracked and segmented human body can be generated and visualized. Figure 3.8 presents the basic skeleton of the segmented human body in the captured depth image, generated and provided by the NiTE middleware. NiTE is able to provide the skeletal representations of up to two human bodies simultaneously. Further, the NiTE middleware includes various calibration and smoothing functions to enhance the skeleton recognition functionality. In addition to the user segmentation and skeleton generation functionality, the NiTE middleware additionally provides pre-defined methods for hand tracking as well as algorithms for pose and gesture recognition.

3.1.4 GLUT

For the graphical output of the HomeRehab application, the *OpenGL Utility Toolkit* [6] (GLUT) library has been used. This library contains a basic application programming interface for graphics programming with OpenGL. The main window and its embedded subwindows are created, managed and constantly refreshed through the toolkit. Within these subwindows, the toolkit

---

[6] GLUT - The OpenGL Utility Toolkit, `http://www.opengl.org/resources/libraries/glut/`, Accessed: 2014-02-11

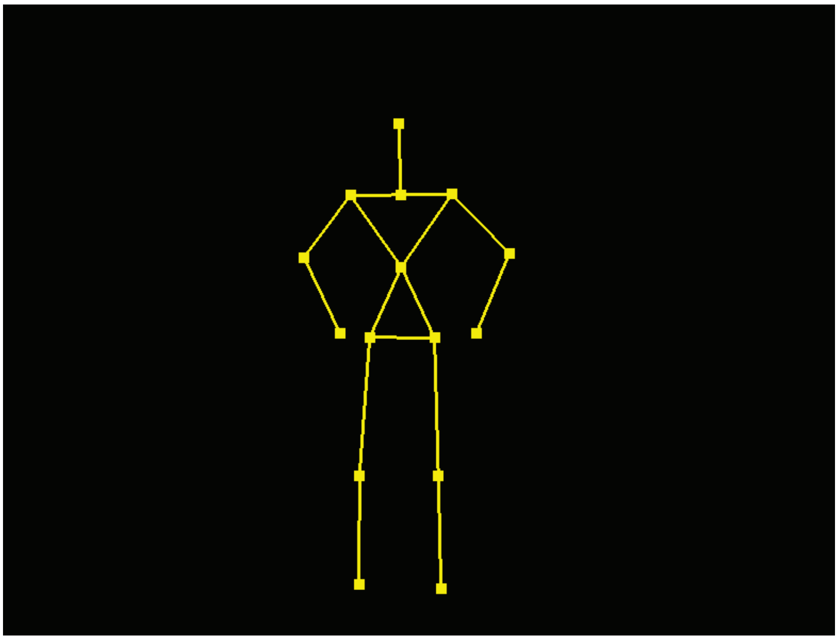

Figure 3.8: Skeletal representation of the segmented user in the depth image, consisting of 15 skeletal joint positions and generated and provided by the NiTE middleware.

manages the visualization of the the captured color information of the connected depth-sensor, as well as of the pre-recorded motion-files loaded. Beside that, skeletal user representation provided by the NiTE middleware is overlaid on the screen through the toolkit. Further, all visual overlays and text notifications within the application windows are generated with this utility toolkit. Additionally, the GLUT library also provides a simple keyboard input listener which is used to process specific input keys required to change several settings within the application.

3.2 Implementation

In this section the actual implementation of the HomeRehab application is presented. Initially, the hardware components and software tools used for development and implementation are listed. After that, the basic design of the application interface including the automated scaling functionality is introduced. Subsequently, the single components the application comprises of as well as their interplay are examined and the basic structure of the application is shown. Beside that, the temporal sequence of the application is presented and the different program states are introduced. And finally the implemented motion analysis and frame compensation algorithms are examined in detail.

3.2.1 Hard- and Software

The hardware components used for development, implementation and evaluation of this project are limited to a conventional Laptop-computer and a low-cost depth-sensing camera:

- ASUS Zenbook UX31A
  - Intel Core i7 3517U Processor
  - Windows 7 Professional 64 bit
  - 4GB DDR3 RAM
- ASUS Xtion PRO LIVE
  - With firmware update 5.8.22
  - Connected via USB 3.0

The software tools used are an IDE for implementation and three open-source libraries:

- Microsoft Visual Studio Professional 2012
  - Programming language used for implementation: C++
- OpenNi 2.2 64 bit
- NiTE 2.2 64 bit
- GLUT - OpenGL Utility Toolkit

3.2.2 User Interface Design

The user interface of the HomeRehab application is kept very simple and consists of a main window and three embedded subwindows. The main focus lies on Subwindow 1 and Subwindow 2 which are arranged side-by-side and occupy the largest space of the application window. The therapist's motions while performing the reference rehabilitation exercises provided by motion-files are shown in Subwindow 1 while Subwindow 2 shows the movements of the patient though the data provided by the connected depth-sensing camera. The output in Subwindow 1 is generated by the *Reference-Viewer* component of the application. It also includes an overlay showing the current playback speed of the currently loaded reference rehabilitation exercise. Subwindow 2 is controlled by the *Live-Viewer* component. It includes an overlay showing the patient's reached scores representing the exercise execution accuracy. The single application components are described in detail in the following section. A third subwindow beneath the other two subwindows is used to output debugging information witch has been used while implementing and testing the prototype application. Figure 3.9 visualizes the basic user interface design of the HomeRehab application and the actual implementation of the prototype interface.

The size of the main window is adjustable at will. By simply defining the desired width of the application window, the sizes of all subwindows are adjusted automatically while the aspect ratio of the data provided by the depth-sensing camera is taken into consideration. The resolution

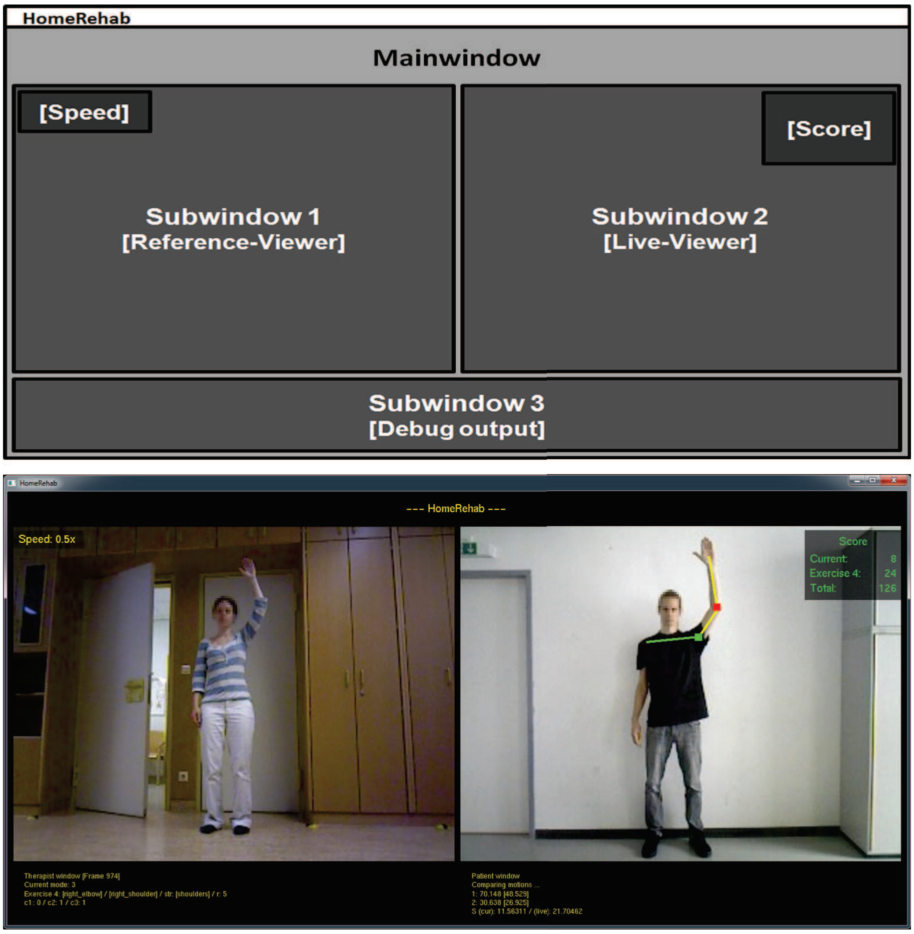

```
Therapist window [Frame 974]
Current mode: 3
Exercise 4: [right_elbow] / [right_shoulder] / str: [shoulders] / r: 5
c1: 0 / c2: 1 / c3: 1
```

```
Patient window
Comparing motions ...
1: 70.148 [48.529]
2: 30.638 [26.925]
S (cur): 11.56311 / (live): 21.70462
```

Figure 3.9: User interface design of the HomeRehab application consisting of a main window and three embedded subwindows and its implementation in the prototype application.

of the data provided by the depth-sensing camera used in this project is 640x480 pixels resulting in an aspect ratio of 1.333. To avoid distortions in the visualized motion data, the aspect ratios of Subwindow 1 and 2 are constantly kept on this value, no matter what application window width has been chosen.
For a better intelligibility and readability, Subwindow 1 will be referred to as the *therapist-screen* for the rest of this thesis, while Subwindow 2 will be referred to as the *patient-screen*.

3.2.3 Basic Application Design

For a better understanding of the functionality of the whole application, comprising of 2900 lines of code in 7 classes written in C++, and for a better overview of the whole program structure, the application can be divided into six interdependent components as shown in Figure 3.10.

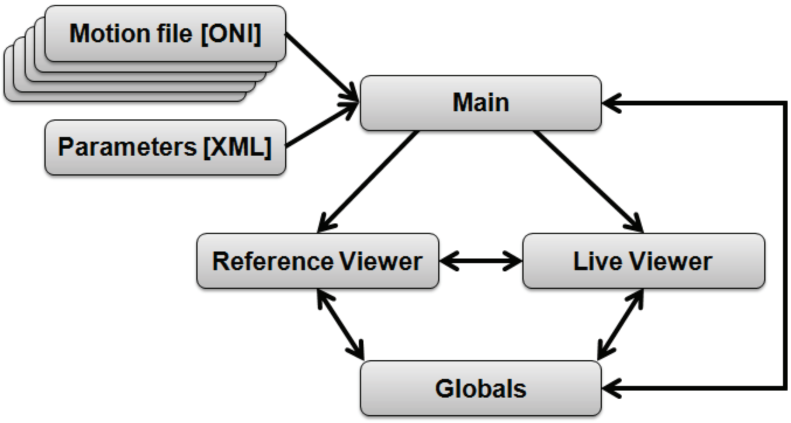

Figure 3.10: Basic structure of the HomeRehab application implementation.

In the following subsections, the tasks and functionalities of the single components, as well as the interaction between the components, are described.

Motion-Files

The first application component comprises of the pre-recorded **motion-files**, each containing one single and non-recurring rehabilitation exercise. These files can be generated through the *NiViewer* application which is freely available and provided by OpenNI. The NiViewer application includes a recording function which is able to store the data of a connected depth-sensing camera into specific motion-files. With this application in combination with a depth-sensing camera, a therapist can provide the needed motion-files through recording his or her motions while performing a correct rehabilitation exercise. These motion-files are then used by the HomeRehab application as a reference for correct exercise execution. The single files are named after their containing exercise followed by the OpenNI file format ending *.ONI* (e.g. *shoulder_abduction_exercise.ONI*) and include data from the device's RGB-camera as well as from the depth-sensor.

XML-File

In addition to the motion files, the therapist also provides one single **XML-file** (Extensible Markup Language) including exercise parameters which represents the second component of the HomeRehab application. This XML-file contains all necessary additional exercise information and all adjustable parameters regarding every single rehabilitation exercise provided in the motion-files. In this XML-file, every rehabilitation exercise the respective patient has to perform at home is listed with an exercise name and an unique identification number. Beside that, the file provides the filename of every single motion-file containing the exercises, as well as the respective exercise start- and end-frames. However, the main issue of this XML-file is to provide exercise related parameters that can be individually adjusted for each patient, considering his special needs. With these parameters, the following attributes can be defined separately for every exercise:

- Exercise criteria
- Tolerance
- Number of repeats
- Exercise description

These adjustable parameters in the XML-file are explained in detail in the following subsections.

Exercise Criteria The first adjustable parameters provided in the XML-file define the criteria for a correct exercise execution. For every exercise, one to three criteria can be chosen that have to be met by the patient at any time while performing this exercise. Only if all of the defined criteria for an exercise are met steadily while exercise execution, this exercise is classified by the application as performed correctly. If one or more criteria are not met at any time of the exercise execution, the exercise is classified as performed incorrectly. Two different types of exercise criteria can be set.

The first type of selectable exercise criteria are angles between limbs that have to remain within specific limits while performing the respective exercise. By defining a specific angle as an exercise criterion, the application calculates and tracks this angle in the motions performed by the therapist showing the rehabilitation exercise, as well as in the motions performed by the patient while mimicking this exercise, and compares them respectively. Thus, by setting angles as exercise criteria, the therapist can define which parts of the body are addressed by the exercise and have to be considered for a correct exercise execution. For instance, at an exercise where the patient has to perform specific motions with the right arm, the criteria for this exercise would be set to 'right_elbow' and 'right_shoulder' in the XML-file, meaning that the application validates if the patient's right elbow and right shoulder angles remain to stay within specific limits while executing the exercise. The following eight different angles can be set as exercise criteria:

- Right or left shoulder angle
- Right or left elbow angle
- Right or left hip angle
- Right or left knee angle

The second type of selectable exercise criteria regards the patients' body posture, where typical rehabilitation exercise requirements, like keeping the shoulders or hips straight while performing the exercise, can be defined. By defining this type of criteria for an exercise, the application validates if the patient is fulfilling specific posture requirements while performing the respective exercise. For instance, the therapist can set the criterion to keep the shoulders straight while performing the exercise. Through this, the application continuously tracks and calculates the difference in height of the patient's right and left shoulder and validates if the patient is raising a shoulder too much while executing the exercise. The following four different posture criteria can be set for each exercise:

- Keep shoulders straight
- Keep hips straight
- Keep the right knee straight
- Keep the left knee straight

In the implemented prototype of the HomeRehab application, a maximum of two criteria regarding angles between limbs and one criterion regarding posture per exercise are supported. At least one criterion has to be defined per exercise in order to enable the application to distinguish between correct and incorrect exercise executions, the other two are optional and can be left blank if not needed.

Tolerance Another adjustable parameter in the XML-file is a tolerance value that can be set for each exercise separately. This parameter specifies a percentage value defining the maximum of the allowed aberrations between the therapists' and the patients' motions while performing the same exercise, considering the defined exercise criteria. Thus, through adjusting the tolerance value for an exercise, the therapist can define how accurately the patient has to perform the required exercise for classifying the exercise execution as correct. Through the tolerance attribute, the application calculates the angle and posture limits representing the border between correct and incorrect exercise executions. If the patients' relevant angles or body postures are beyond these limits at any time while performing the exercise, the respective exercise criterion is not met and the exercise execution is classified as incorrect by the application. The calculation of these limits is based on maximum values which have been defined in association with a medical professional and represent a completely wrong exercise execution. The maximum value for aberrations of angles between the therapist and the patient is set to 90 degrees. The maximum value for the differences of joint heights for the criteria regarding the straightness of body parts is determined through calculating the angle between the body part to be kept straight and the horizontal plane. The maximum difference in height lies where this angle is 45 degrees. At the criteria to keep the knees straight the angle is calculated in regard to the vertical plane respectively. With these maximum values and with the tolerance value set for each exercise, the limits for the allowed aberrations of the patients motions from the reference motions can be calculated for each exercise. For a better understanding, examples with a tolerance value set to 25% are given in the following.

Assuming that an exercise criterion is the angle of the patient's right elbow, a tolerance value of 25% means that the patient is allowed to differ the elbow posture to up to 25% of the maximum aberration of 90 degrees compared to the elbow posture of the therapist. Thus, the difference between the elbow angle of the therapist and the patient has to remain within 22.5 degrees at any time while executing the exercise. Assuming that an exercise criterion it to keep the shoulders

straight while performing the exercise, a tolerance value of 25% means that the height difference between the patient's left and right shoulder has to stay beneath the height difference of the therapist's right and left shoulder plus 25% of the calculated maximum difference.
With the help of this tolerance parameter, every exercise can be adjusted to the current state of health and rehabilitation progress of the patient. As the patient's recovery progresses, the tolerance value can be reduced continuously to keep the exercises challenging and boost further recovery.

Number of Repeats The next adjustable parameter in the XML-file containing the exercise parameters defines the number of required repeats for each exercise, as some exercise have to be repeated more often than others. This value is enabling the application to repeat the respective exercise as often as the therapist suggests before loading the next exercise.

Exercise Description The final adjustable parameter is a textual description of the exercise which can be defined for each exercise in the XML-file, providing the patient with instructions regarding the exercise execution which are shown on the screen while the exercise is being loaded. This parameter is optional and can be left blank if not needed.

An example excerpt of a XML-file containing specific exercise parameters is shown in Figure 3.11.

```
 ⋮
<exercise>
  <number>4</number>
  <name>Shoulder joint - Lift arm</name>
  <filename>lift_right_arm.oni</filename>
  <exerciseStartFrame>490</exerciseStartFrame>
  <exerciseEndFrame>1070</exerciseEndFrame>
  <relevantAngle1>right_elbow</relevantAngle1>
  <relevantAngle2>right_shoulder</relevantAngle2>
  <keepStraight>shoulders</keepStraight>
  <tolerance>20</tolerance>
  <numberOfRepeats>5</numberOfRepeats>
  <description1>- Keep knees slightly bent, tense your stomach, tense your pelvic floor</description1>
  <description2>- Keep right arm slightly bent, thumb points up</description2>
  <description3>- Raise arm above your head on the side of your body, keep shoulder down</description3>
  <description4>- Lower your arm</description4>
</exercise>
<exercise>
  <number>5</number>
  <name>Cervical spine rotation</name>
  <filename>cervical_spine_rotation.oni</filename>
 ⋮
```

Figure 3.11: Example excerpt of a XML-file specifying exercise parameters.

Main-Class

The third application component in the HomeRehab application is the **Main-class** which represents the entry point of the software. This class can be seen as the central component of the whole application and interacts with all other application components. The Main-class has several different tasks to fulfill:
Initially, the *OpenGL Utility Toolkit* is loaded and initialized through witch the application main window is generated and opened. After that, all necessary subwindows are created and aligned within this main window. Once all windows are set up, the next task of the Main-class is to load and analyze the XML-file with the exercise parameters provided by the therapist. For this step a customized XML-parser which is considering the fixed structure of such a XML-file containing exercises parameters, has been implemented. This parser reads the necessary information and parameters regarding the exercises line by line from the provided XML-file and transfers them to the internal storage, making the data available to all application components at any time. After all parameters are extracted from the file, the Main-class is preparing the communication with the connected depth-sensing camera. The OpenNI-framework is being initialized and the data-streams of both the RGB-camera as well as of the depth-sensor of the connected device are created and started. Once the OpenNI framework is fully set up, the NiTE middleware is loaded and initialized for activating the user segmentation and skeleton tracking functionality. Another component of the Main-class is the implemented keyboard listener which is steadily checking for specific keyboard inputs throughout the application execution. As soon as pre-defined keys are pressed, corresponding events are triggered within this class. The Main-class is also responsible for constantly refreshing the output of the main- and subwindows if new data is available and needs to be drawn. And finally the Main-class creates, initializes and controls the Reference- and the Live-Viewer classes which are representing the next two application components.

Reference-Viewer

The **Reference-Viewer** is the fourth component of the HomeRehab application. At the same time, this is the component with the most functionality and complexity. As the name suggests, the Reference-Viewer is responsible for all tasks regarding the reference exercises provided by the therapist through the motion-files. The three main functionalities of the Reference-Viewer component are loading the pre-recorded motion-files containing the rehabilitation exercises to be performed, analyzing the therapist's motions stored in the files and visualizing the exercises on the screen for the patient.
Initially, the first exercise transferred form the provided XML-file containing the exercise parameters to the storage is checked for the filename of the corresponding motion-file. In the next step, the exercise motion-file is loaded and the data-streams of the RGB-camera and the depth-sensor provided in the file are created and started. For analyzing the motions of the therapist, the user-tracker of the NiTE middleware is started, enabling user segmentation and skeleton tracking in the raw depth-data streamed from the motion-file.
After that, the Reference-Viewer class switches through four different program states before loading the next exercise and repeating the same steps again until every exercise extracted from the XML-file is processed. These four program states are visualized in Figure 3.12 and described

in the following. Through these program states, a temporal sequence for the whole application is generated which is repeated with each exercise loaded.

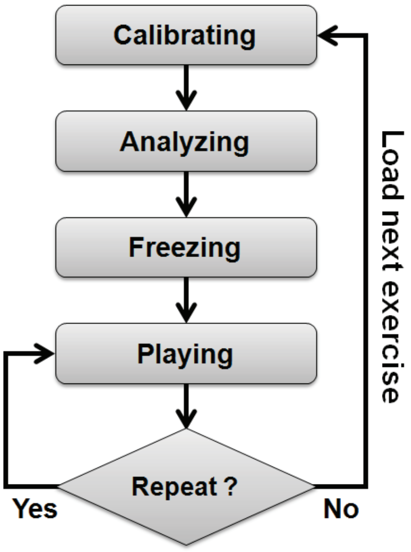

Figure 3.12: The four program states the HomeRehab application switches through for every loaded rehabilitation exercise.

Calibrating-Mode The first program state executed after a new exercise has been loaded is the *Calibrating-mode*. In this mode, two tasks are being performed simultaneously. The currently loaded motion-file is streamed frame by frame in the background until the user segmentation is completed and the therapist's skeleton is tracked. While the user segmentation and skeleton tracking processes are being set up, the patient is provided with a loading screen presenting textual information about the currently loaded exercise, which is shown on the therapist-screen. Beside the exercise name and the unique number of the exercise, a textual description about the correct exercise execution is displayed on the screen. The patient is also informed about the required number of repeats of the current exercise. Again, all textual information to be displayed for every exercise can be defined by the therapist through the XML-file. As soon as the defined exercise-start-frame of the motion-file streamed in the background is reached, the application checks whether the therapist's skeleton is already tracked properly. If the skeleton is not yet tracked properly when the exercise-start-frame is reached, the loaded motion file is set back to the first frame and the user segmentation and skeleton tracking procedures are repeated while the loading screen is still displayed. When the defined exercise-start-frame is reached and the therapist's skeleton is tracked properly, the application automatically switches to the next program state which is called the *Analyzing-mode*.

Analyzing-Mode The purpose of this program state is twofold. On the one hand, the exercise motion-file is played back on the therapist-screen, showing the patient the correct exercise execution. While the exercise is played back, a text message pops up on the patient-screen, notifying the patient to watch the exercise execution on the therapist-screen attentively. On the other hand, the second task in this program state is, as the name suggests, to analyze the loaded exercise motion-file and store the relevant reference motions of the therapist. The definition which motions are relevant for the currently loaded exercise is determined through the criteria set in the XML-file. Every criterion is visualized as an overlay on the therapist-screen while playing back the rehabilitation exercise, to additionally inform the patient which parts of the body the shown exercise addresses. If the criteria include specific angles, the angles between the respective limbs of the therapist are calculated and stored for every frame of the loaded motion-file. Same is the case for the criteria regarding keeping the shoulders or hips straight, where the height difference of the left and right shoulder or the left and right hip joints is calculated and stored. At criteria regarding keeping specific body parts vertically straight, the laterally joint differences are stored respectively.
As soon as the end of the exercise shown on the therapist-screen is reached, the application automatically switches to the *Freezing-mode*.

Freezing-Mode In the Freezing-mode, the therapist-screen shows the initial exercise position and freezes the image while and the patient again is notified through a text message appearing on the patient-screen to get into the shown position and prepare for the beginning of the exercise. After this text message disappears, a countdown-timer is shown on the patient-screen counting down from six seconds to zero, preparing the patient for the beginning of the exercise.
As the countdown-timer depletes, the final program state for the loaded exercise is started which is called the *Playing-mode*.

Playing-Mode In this stage of the program, the loaded exercise is played back continuously on the therapist-screen until the defined number of repeats for this exercise is reached. The patient simultaneously performs the shown rehabilitation exercise while being provided with live visual feedback, overlaid on the patient-screen, regarding the patient's exercise execution. After reaching the defined number of repeats, the application stops the data streams of the loaded motion-file, unloads the file and automatically loads the next exercise listed in the XML-file. The loading screen with the textual information about the next exercise is shown again and all four program states are switched through again.

Live-Viewer

The fifth component of the application is the **Live-Viewer**, which is executed simultaneously to the Reference-Viewer component. The Live-Viewer component is responsible for the communication with the connected depth-sensing camera and thus for tracking and analyzing the patient's motions. It also manages the visual output on the patient-screen, showing the RGB-data stream of the connected device as well as the reached scores and the visual feedback regarding the patient's exercise execution. The data of the RGB-camera of the connected device is streamed and shown on the patient-screen continuously from application start to end, meaning that the patient always sees the own motions in this subwindow. While the program is within the *Playing-mode*, where the therapist's exercise execution is shown on the therapist-screen while the patient performing the same exercise simultaneously is shown on the patient-screen, the motions of the patient are analyzed regarding the criteria defined for the currently loaded exercise. In this mode, the task of the Live-Viewer component is to compare the patient's motions to the reference motions of the therapist provided in the loaded motion-file and to calculate aberrations respectively. Based on these aberrations live visual feedback is generated and provided as an overlay on the screen showing the patient's movements. These aberrations are triggering the visual feedback and are further the basis for calculating the scores representing the patient's exercise execution accuracy. The scores are overlaid on the patient-screen while the *Playing-mode* is active and are updated after each single exercise cycle. Beside that, the Live-Viewer component is also responsible for generating and visualizing the text messages popping up on the patient-screen, which are guiding the patient through the exercises. Another functionality implemented in this application component is the timer-function which is responsible for generating a constant and smooth temporal sequence while processing every single rehabilitation exercise.

Globals-Class

The sixth and last component of the HomeRehab application is the **Globals-class**. The main task of this component is the information exchange between the other components. It contains all global constants and variables as well as methods that are needed in more than one other application component. This is also the place where the OpenNI-, NiTE- and GLUT-libraries are embedded into the application enabling all other application components to access the functionality of these external libraries by simply including the Globals header file.
In the following subsections the motion analysis and frame compensation algorithms implemented in the application will be examined in detail.

3.2.4 Motion Analysis

The actual motion analysis in the HomeRehab application takes place while the program is within the *Analyzing-mode*, where the motions of the therapist are captured through the data of the motion-files loaded, and while the program is within the *Playing-mode* where the motions of the patient are captured through the data provided by the connected depth-sensor. The motion analysis of the application is depending on the defined criteria in the parameters of each exercise. If exercise criteria regarding angles are defined, the 3D-positions of the respective skeletal joints are examined through the NiTE middleware and the angles between these joints are calculated separately for each frame. The limbs of the tracked skeleton, represented by connected joint positions, are treated as 3-dimensional vectors. The cosine angle between two 3-dimensional vectors can be calculated through the dot product of the two vectors divided by the product of their scalar values and converted to degrees.
Equations 3.1 and 3.2 show this angle calculation in the example of shoulder and elbow angles.

$$ShoulderAngle = arcos\left(\frac{arm \cdot trunk}{||arm||\,||trunk||}\right)\left(\frac{180}{\pi}\right) \tag{3.1}$$

$$ElbowAngle = arcos\left(\frac{forearm \cdot arm}{||forearm||\,||arm||}\right)\left(\frac{180}{\pi}\right) \tag{3.2}$$

Angles between two 3-dimensional vectors calculated through these equations are invariant to translation, rotation and scale. Regarding the HomeRehab project, this invariance means that the calculated angles are independent from the patient's proportions as well as from the distance of the patient to the depth-sensor. Also the orientation of the patient's body to the sensor does not influence the calculated angles as long as the skeleton tracking is not interfered through occlusions. In the *Analyzing-mode* the relevant angles of the loaded motion-file containing the therapist's exercise execution, calculated through these equations, are stored into an array for every streamed frame. In the *Playing-mode*, where the patient is mimicking the shown exercise, the relevant angles of the patient's limbs are calculated for every streamed frame through the data of the depth-sensor tracking the patient's movements and are compared to the reference angles stored in the array while analyzing the therapist's motions. The application checks for every streamed frame of the connected depth-sensor to which amount the current angles calculated in the patient's skeletal representation differ from the reference angles stored in the array for the respective frame. If the difference of an angle is exceeding the specified tolerance value for this exercise at any time, the application marks the respective angle and validates if the difference is exceeding the limit for more than one second. If this is the case, this exercise criterion is counted as not met and the patient's exercise execution is classified as incorrect by the application. When the patient corrects the posture within one second, the criterion is still counted as met. Through this approach, short deviations between the calculated joint angles of the patient and the stored joint angles of the therapist are allowed.

For exercise criteria regarding the task of keeping specific body parts straight, a different motions analysis approach is used. In this approach the respective skeletal joints are analyzed in regard of their aberration to the horizontal or vertical axe. In the example of the criterion to keep the shoulders straight this means that the difference of height of one shoulder in regard to the other shoulder is calculated. When keeping the shoulders perfectly straight this difference would be zero. The values of this difference occurring in the reference exercise of the therapist are stored while analyzing the therapist's motions in the *Analyzing-mode*. In the *Playing-mode*, the difference values occurring in the patient's movements are calculated for every frame and compared to the stored values. If at any time while performing the exercise the difference values in the patient's skeleton exceed those in the therapist's skeleton more than the threshold defined by the set tolerance value allows, the straightness criterion is not met and the exercise execution is classified as incorrect by the application.

3.2.5 Compensating Lost Frames

Due to the fact that the HomeRehab application is acquiring data from the loaded motion-files and from the connected depth-sensing camera simultaneously, single frames happens to be skipped and through this their information is lost. The reason for this is estimated in too low performance of the personal computer used for testing the system. To be independent from the computer hardware used, a frame compensation algorithm has been implemented to deal with these lost frames. Since the OpenNI framework assigns an unique consecutive number to every streamed frame it is possible to constantly check if a frame is missing. While analyzing the therapist's motions in the provided motion-file, the lost frames are detected by the application and compensated through the data of the previous frame and the data of the next frame. For instance when analyzing and saving angles between limbs, the angle in the lost frame is determined through the mean value of the previously calculated angle and the next calculated angle. Through this approach the stored data after analyzing an exercise motion-file contains only valid values for every frame of the respective exercise. Otherwise a valid comparison of the patient's motions to the motions of the therapist for every frame would not be possible and the application would provide false-negatives when classifying whether the patient's exercise execution is correct or not.

3.3 The HomeRehab-Application

In this section, the final prototype application is presented, application screenshots are shown and the functionalities and temporal sequence of the application are examined. A screenshot of the final prototype application interface, while the application is loading the first rehabilitation exercise is shown in Figure 3.13. On the left side (therapist-screen) the exercise loading screen presenting textual information about the current exercise is shown. Beside the unique exercise number and the name of the exercise, a textual description of the correct exercise execution is presented and the patient is additionally informed about the number of required repeats for this exercise. The screen on the right side (patient-screen) visualizes the data of the connected depth-sensing camera and shows the patient's movements. The patient sees himself mirrored within this subwindow which means that his motions are shown on the screen the same way as if he would stand in front of a mirror which is, regarding to Dr. Paternostro-Sluga, common for performing rehabilitation exercises at home. The user segmentation and skeleton tracking procedures are executed initially after the application start. As long as the skeleton of the patient is not yet tracked properly, which usually takes not longer than five seconds, a red text message appears on the patient-screen which can be seen in the figure as well. As soon as the patient's skeleton is tracked properly, this text overlay disappears. If the depth-sensor loses the tracking of the patient at any time while executing the application, the text overlay is shown again until the skeleton of the patient is tracked properly again. Alternatively the text overlay can be set to be visible at any time to visualize the current tracking status throughout the application execution by pressing the *l-key* on the keyboard. On the bottom of the application window (Subwindow 3), textual information about the current program state as well as about various implementation related values is shown, which has been used for debugging while implementing the application.

As soon as the user segmentation in the data of the loaded motion-file is completed, the skeleton of the therapist is tracked properly and the start-frame of the loaded exercise is reached, the application automatically switches to the next program state which is the Analyzing-mode. The loading screen shown on the therapist-screen, providing the textual information about the current exercise disappears and the exercise execution of the therapist, taken from the loaded motion-file, is shown. A screenshot of the application in this program state is shown in Figure

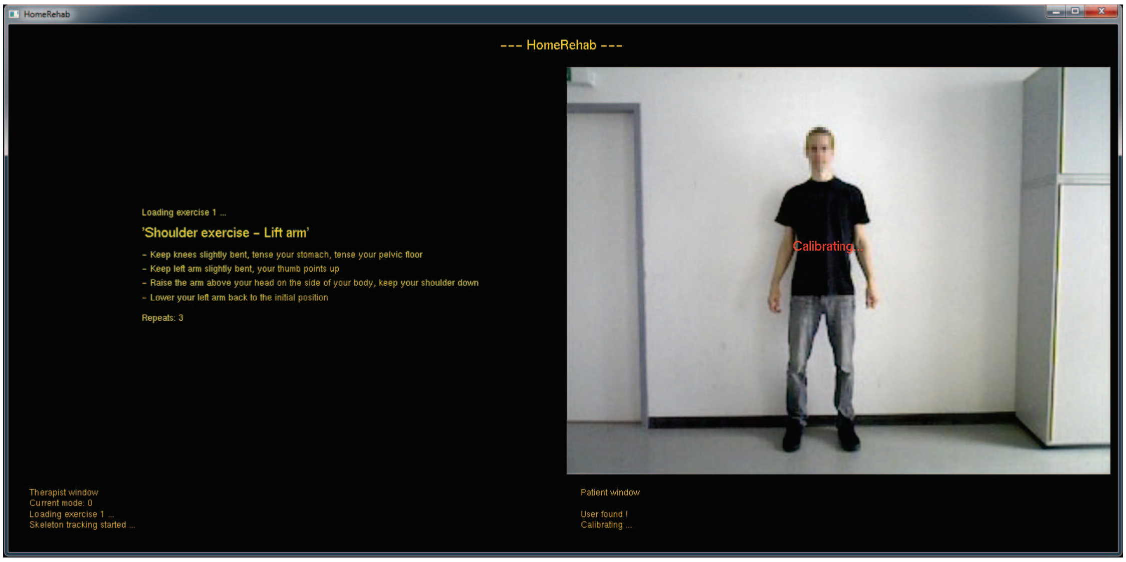

Loading exercise 1 ...

'Shoulder exercise – Lift arm'

- Keep knees slightly bent, tense your stomach, tense your pelvic floor
- Keep left arm slightly bent, your thumb points up
- Raise the arm above your head on the side of your body, keep your shoulder down
- Lower your left arm back to the initial position

Repeats: 3

Therapist window
Current mode: 0
Loading exercise 1 ...
Skeleton tracking started ...

Patient window

User found !
Calibrating ...

Figure 3.13: Screenshot of the HomeRehab-application within the Calibrating-mode while loading the first rehabilitation exercise. A loading screen, providing the patient with textual information about the current exercise, is shown on the therapist-screen. At the same time, the patient's motions are shown on the patient-screen through visualizing the color information of the connected depth-sensor tracking the patient. Additionally, the patient is informed that the user segmentation and skeleton tracking processes are not yet completed, through a text notification overlaid on the patient-screen.

3.14. While the therapist is shown, a text notification on the patient-screen reminds the patient to watch the exercise attentively. Beside the correct exercise execution itself, the patient also sees which parts of the body are addressed by this rehabilitation exercise and which criteria have to be met while performing this exercise. This information is shown through visual elements on the therapist-screen, overlaid on the therapist's body. In the example exercise shown in the figure, the criteria for correct exercise execution are two angles (left shoulder angle and left elbow angle) as well as the criterion to keep the shoulders straight. This information is shown in the form of two purple dots, overlaid on the therapist's shoulder and elbow and in the form of a straight purple line connecting the therapist's shoulders. While the therapist's exercise execution is shown, the application analyzes and stores the required motion information.

When the end of the rehabilitation exercise shown by the therapist is reached, the application freezes the therapist-screen while showing the therapist in the initial position of the exercise. The application has now reached the Freezing-mode where the patient is asked to get into the shown position through a text notification popping up on the patient-screen. A screenshot of the application in this program state is shown in Figure 3.15.
Subsequently, this text notification on the patient-screen disappears and a countdown timer, counting down from six seconds to zero is instead shown on this screen. Within these six sec-

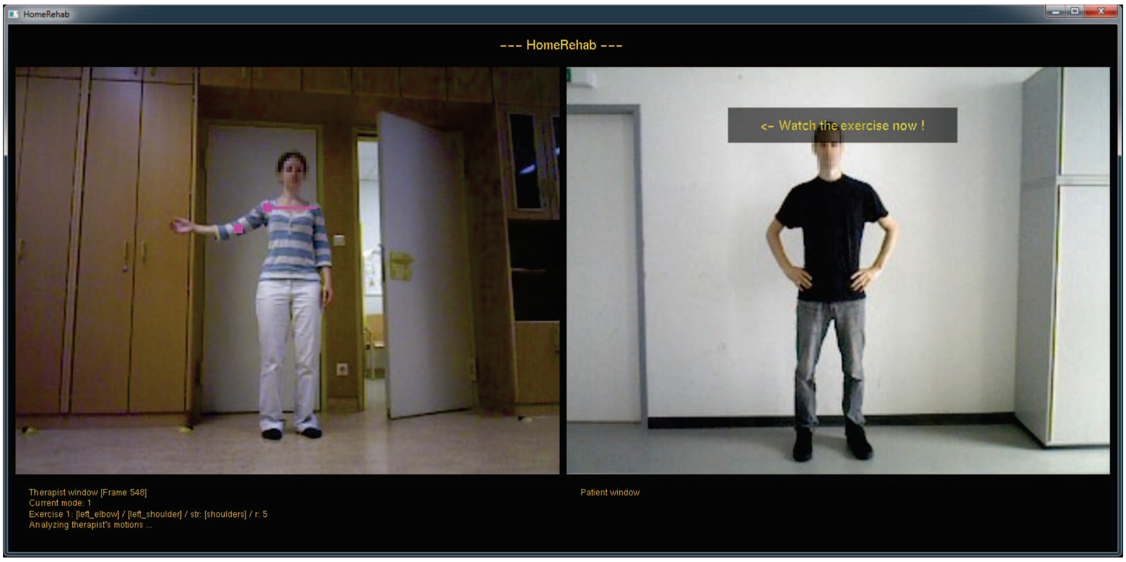

<- Watch the exercise now !

```
Therapist window [Frame 548]
Current mode: 1
Exercise 1: [left_elbow] / [left_shoulder] / str: [shoulders] / r: 5
Analyzing therapist's motions ...
```

Figure 3.14: Screenshot of the HomeRehab-application within the Analyzing-mode. On the therapist-screen the correct exercise execution performed by the therapist is shown to the patient. Through overlays on the therapist's body, the patient is additionally informed about the relevant criteria for a correct exercise execution. At the same time, the relevant motions of the therapist are analyzed and stored by the application. The patient is asked to watch the exercise through a text message popping up on the patient-screen.

onds, the patient can prepare and get ready for the exercise begin.

As soon as the countdown timer depletes, the application is automatically switching to the last program state for this exercise (Playing-mode) where the actual rehabilitation exercise execution of the patient takes place. The image on the therapist-screen stops freezing and the shown therapist starts with the performance of the correct exercise execution while the patient mimics the therapist's motions simultaneously. A screenshot of the application in this program state is shown in Figure 3.16. The therapist's exercise execution is shown on the therapist-screen, while the patient-screen shows the patient while mimicking the therapist's motions. As shown in the figure, this program state includes three additional visual elements overlaid on the screens. On the therapist-screen, an additional information about the current playback speed is shown, while on the patient-screen the visual feedback regarding the patient's exercise execution as well as the reached scores of the patient are overlaid. These visual elements are described in more detail in the following subsections.

3.3.1 Playback Speed

While the application is within the Playing-mode where the patient is executing the rehabilitation exercises simultaneously to the therapist shown, the playback speed of the currently loaded motion-file and respectively the exercise execution speed of the therapist can be adjusted.

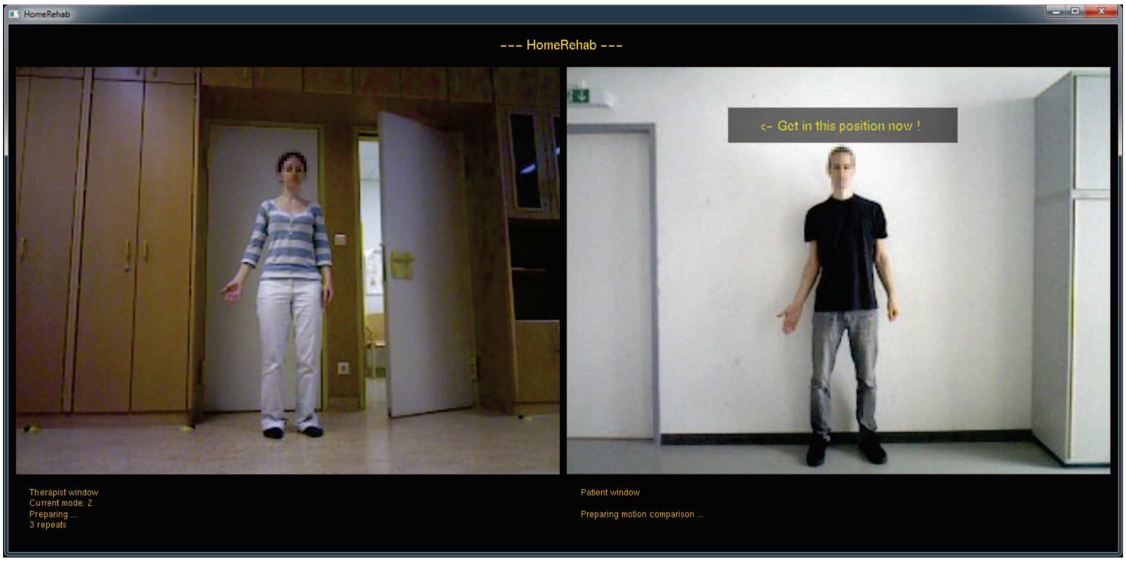

--- HomeRehab ---

<- Get in this position now !

Therapist window
Current mode: 2
Preparing ...
3 repeats

Patient window

Preparing motion comparison ...

Figure 3.15: Screenshot of the HomeRehab-application within the Freezing-mode. The output on the therapist-screen freezes while showing the therapist in the initial position of the exercise. The patient is asked to get into the same initial exercise position through a text notification popping up on the patient-screen.

Through the *Up-Arrow* and *Down-Arrow* keys of the keyboard, the playback speed can be raised or lowered by the patient at any time if the therapist's movements are too slow or too fast for the patient's current condition. The current playback speed is represented by a multiplier shown in the left upper corner of the therapist-screen. A multiplier means that the playback speed is currently multiplied by the shown value. The default value for every exercise is 1.0, resulting in a playback speed of 30 frames per second. A value of 1.0 results in playing back the exercise execution at the original speed, meaning that the therapist performed the exercise while recording the motion-file at the exact same speed as the exercise is played back by the application. Every key press raises or lowers the current playback speed one-tenth, until a maximum playback speed of 60 frames per second or a minimum playback speed of 3 frames per second is reached. Through this, the patient can precisely adjust the speed of the exercise execution to the current needs. The selected playback speed multiplier is kept for every repetition of the currently loaded exercise as well as for all exercises loaded afterwards, until the value is adjusted again or the home rehabilitation session ends.

3.3.2 Visual Feedback

While the patient is performing the required rehabilitation exercises, the application provides live and distinct visual feedback about the correctness of the patient's current motions in regard to the motions of the therapist. This feedback is visualized on the patient-screen and overlaid directly on the patient's body. The overlaid feedback elements are color-coded. The color green

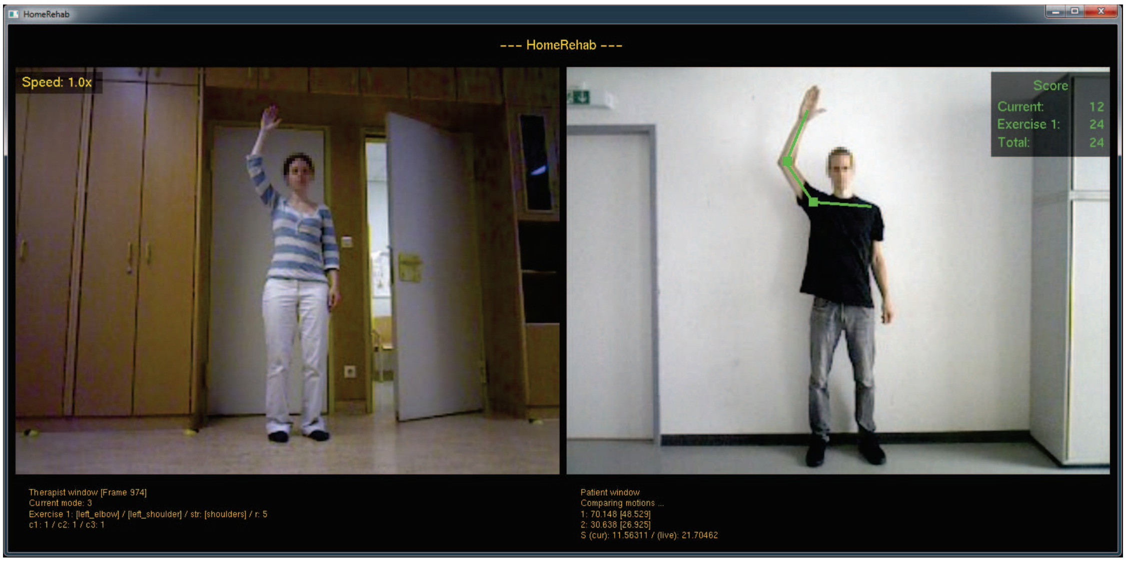

```
Score
Current:     12
Exercise 1:  24
Total:       24
```

```
Therapist window [Frame 974]
Current mode: 3
Exercise 1: [left_elbow] / [left_shoulder] / str: [shoulders] / r: 5
c1: 1 / c2: 1 / c3: 1
```

```
Patient window
Comparing motions ...
1: 70.148 [48.529]
2: 30.638 [26.925]
S (cur): 11.56311 / (live): 21.70462
```

Figure 3.16: Screenshot of the HomeRehab-application within the Playing-mode. The patient is performing the rehabilitation exercise simultaneously to the therapist shown on the therapist-screen. An additional visual element on the upper left corner of the therapist-screen informs the patient about the current playback speed multiplier. The patient-screen shows the patient while mimicking the therapist's motions while the visual feedback regarding the patient's exercise execution is overlaid on the patient's limbs. The feedback in this case is indicating that the patient is currently performing the exercise correctly. Additionally, an overlay with the patient's reached scores is shown on the upper right corner of the patient-screen.

represents a correct exercise execution while the colors red and yellow are notifying the patient about aberrations from the correct motions. Depending on the criteria defined for each exercise through the XML-file containing the exercise parameters, the visual feedback elements can look differently for each exercise. However, the feedback elements are always based the same principles. If an exercise criterion is that a specific joint angle does not exceed a calculated threshold, the visual feedback is generated as follows:
Every relevant angle is represented by an overlaid dot on the patient's body shown on the patient-screen. For instance if the correctness of the patient's right elbow angle is a criterion for a correct exercise execution, a dot is overlaid on the position where the depth-sensor in combination with the NiTE middleware is estimating the patient's right elbow. This dot is shown green if the calculated angle currently remains within the defined limits or red otherwise. Additionally, the limbs forming this angle are overlaid as well. In the example of the elbow angle, the overlaid limbs forming this angle are the patient's upper and lower arm. In the example of the knee angle, the overlaid limbs would be the patient's upper and lower leg. These overlaid limbs follow the same color code as the overlaid dots representing the angles between the limbs. However, there is one exception from this approach:

The exception regards the case when two relevant angles are defined by the same limb. For example the right shoulder angle and the right elbow angle are both calculated through the position of the patient's right upper arm. In this case it is additionally distinguished whether both angles in the patient's skeletal representation are currently correct or incorrect or if one angle is correct while the other one is incorrect. If both angles are currently correct or incorrect, the overlaid limbs forming these angles are shown green or red respectively. However, if one angle is correct while the other one is currently not correct, the respective limbs are visualized yellow to additionally notify the patient that only one of the angles has to be corrected.

If an exercise criterion is to keep a specific body part straight, the visual feedback is represented through an overlaid horizontal or vertical line. This line again is shown green if the calculated threshold is not exceeded and red otherwise.

Figure 3.17 shows the visual feedback overlay while performing an exercise with three exercise criteria. The criteria are the correct angles of the patient's right shoulder and elbow as well as to keep the shoulders as straight as possible. In Figure 3.17(a) the patient is performing the exercise correctly and every defined exercise criterion is met. Figure 3.17(b) shows the patient's exercise performance while the visual feedback is indicating that the patient's shoulder angle is currently not correct while the angle of the patient's elbow is correct. In Figure 3.17(c) the same holds for the patient's elbow angle. And Figure 3.17(d) shows the visual feedback while notifying the patient that the shoulders are currently not kept straight.

Another example showing the visual feedback elements of the HomeRehab application is presented in Figure 3.18. The patient is shown while performing a balance-rehabilitation exercise with the exercise criterion to keep the hips straight while standing on one leg. Figrue 3.18(a) shows the correct exercise execution while in (b) the patient is leaning towards one side too much. Figure 3.19 shows the patient while performing a shoulder exercise. The exercise criterion is to keep the shoulders straight at any time while performing the exercise. In Figure 3.19(a) the correct exercise execution is shown while in (b) the feedback overlay is indicating that the shoulder is lifted too much.

3.3.3 Scores

While performing the required rehabilitation exercises, the application additionally provides the patient with scores that are calculated depending on the patient's exercise execution accuracy. These scores are presented in an overlay on the upper right corner of the patient-screen. Three different score values are calculated and shown, representing the score for the currently performed exercise, the score for the whole exercise cycle with the defined number of repeats and an overall score representing the performance through the whole home-rehabilitation session including all required exercises. The scores output is updated after every completed exercise repetition and reveals if the patient's exercise execution has been classified as correct or as incorrect by the application. The maximum score per exercise repetition is 12 points which can only be reached if every defined exercise criterion is met by the patient at any time while performing the respective exercise. The only exception from this approach are criteria regarding angles between skeletal joints of the patient. These criteria are still classified as met if the calculated angles are beyond the defined tolerance for a period of less than one second. As soon as the aberrations of the angles are beyond the defined tolerance for more than one second, this

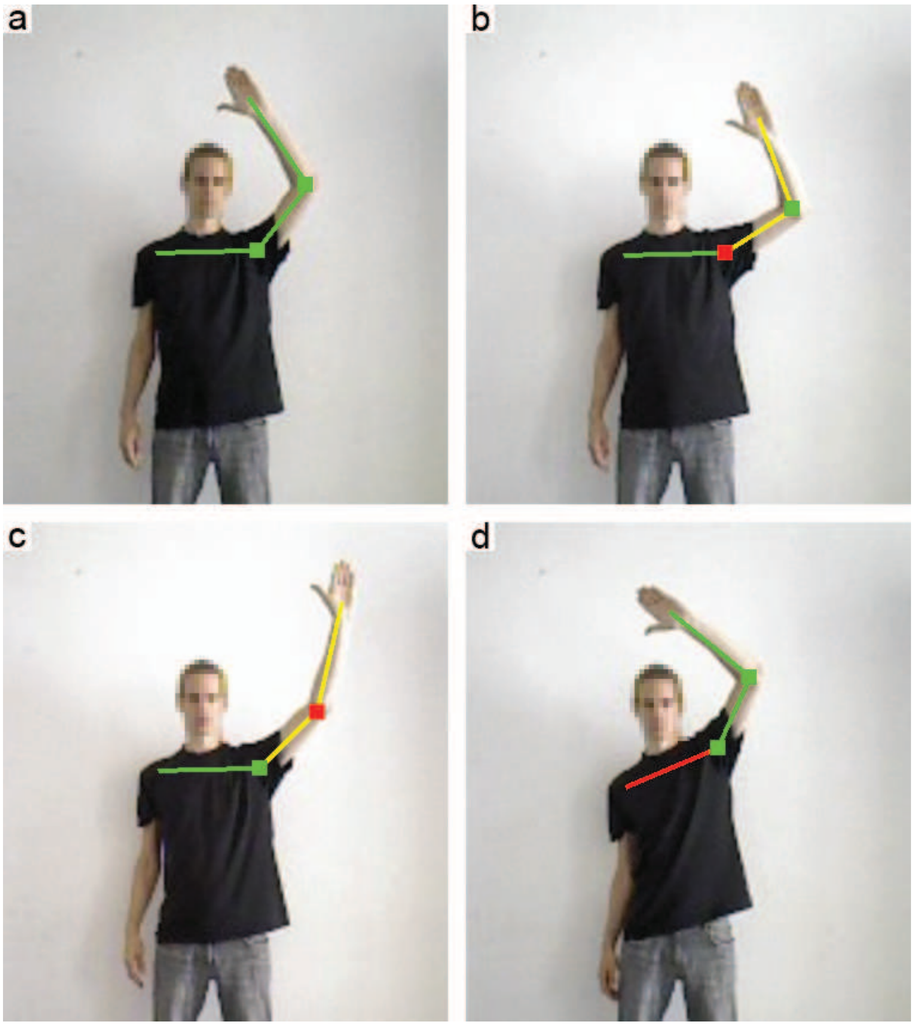

Figure 3.17: Visual feedback elements of the HomeRehab application, overlaid on the patient's body while performing a rehabilitation exercise. The images present the feedback for correct exercise execution (a), aberrations regarding the patient's current shoulder angle (b) and elbow angle (c) and aberrations regarding keeping the patient's shoulders straight.

exercise criterion is classified as not met for the currently performed exercise and the reached score is lowered. Through this approach, short deviations between the calculated joint angles of the patient and the stored joint angles of the therapist are allowed. If one or more exercise criteria are not met by the patient while performing the exercise, the reached score is lowered respectively. At exercises where three criteria are defined by the therapist, every criterion met is worth four points. If only two criteria are defined, each is worth six points, and at exercises with only one criterion to be met, this criterion is worth 12 points. Through this approach, the reached score always indicates how many criteria have been met while executing the exercise in regard

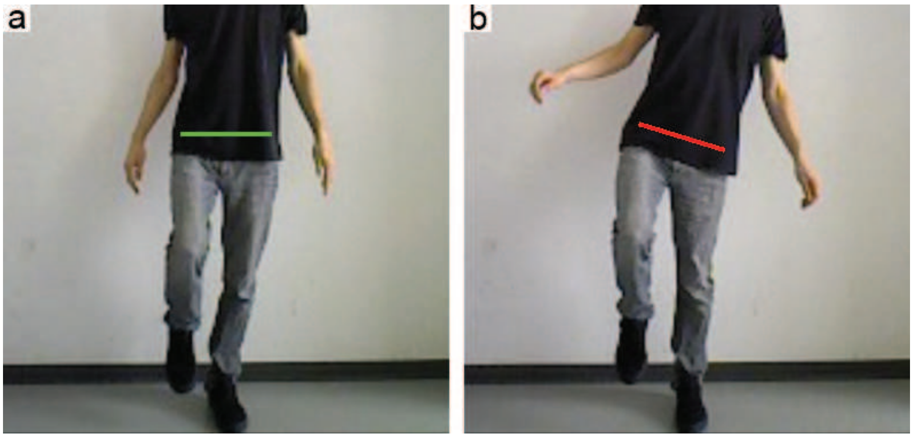

Figure 3.18: Visual feedback of the HomeRehab application regarding the patient's hips, indicating correct (a) and incorrect (b) exercise execution.

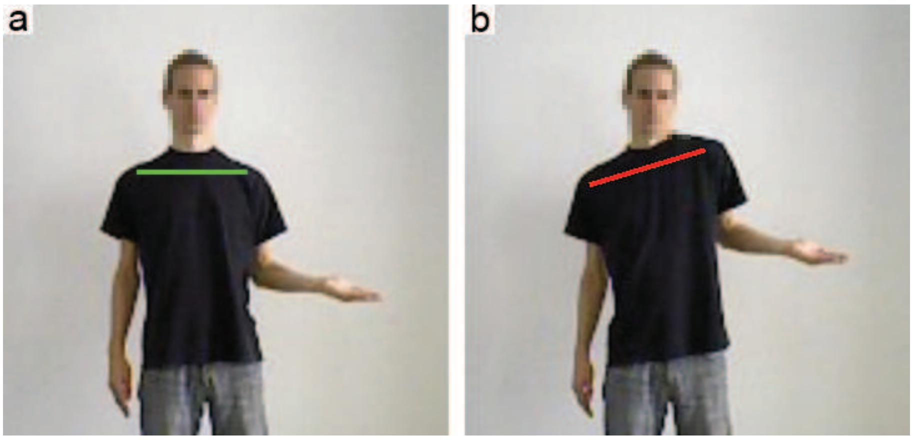

Figure 3.19: Visual feedback of the HomeRehab application regarding the patient's shoulders, showing correct (a) and incorrect (b) exercise execution.

to the number of criteria defined. The following scores can be reached for every repetition of an exercise, indicating the met criteria.

- 12 Points - Every exercise criterion is met
- 8 Points - Two of three defined criteria are met
- 6 Points - One of two defined criteria is met
- 4 Points - One of three defined criteria is met

- 0 Points - No exercise criterion is met

Since a correct exercise execution is only possible if every defined exercise criterion is met at any time while performing the rehabilitation exercise, every score lower than 12 points reached for this exercise indicates that the exercise execution has been classified by the application as performed incorrectly. The purpose of the scoring mechanism is twofold. On the one hand, the overall rehabilitation progress of the patients can be visualized by comparing the recently reached scores of every home-rehabilitation session. On the other hand, the scores are intended to increase the patients' motivation to perform their exercises accurately and to use the application more frequently and thus to get the patients to perform the required rehabilitation exercises more often, resulting in a faster recovery.

As soon as the currently loaded rehabilitation exercise has been repeated as often as defined in the exercise parameters, the application automatically unloads the motion-file and switches back to the Calibrating-mode. The next exercise defined in the XML-file is loaded and the loading screen is shown again, providing the patient with textual information about the new exercise. Subsequently, all program states are switched through again for the new exercise.
Figure 3.20 and Figure 3.21 are presenting the application's therapist- and patient-screen while different rehabilitation exercises are shown by the therapist and performed by the patient.
In Figure 3.20, a balance-rehabilitation exercise is performed by the therapist and mimicked by the patient. The application checks whether the patient's hips are kept straight while rising the leg and provides respective feedback. As can be seen on the therapist-screen, the playback speed and thus the exercise execution speed for this exercise has been lowered by the patient to the half of the original speed.
Figure 3.21 shows the therapist and the patient while performing a neck-related rehabilitation exercise. As the feedback on the patient-screen indicates, the patient is performing the exercise wrong by raising the shoulder too much for a correct exercise execution. Due to the fact that keeping the shoulders straight is the only exercise criterion in this example, and this criterion is not met by the patient, the reached score for this exercise is zero.

Figure 3.20: Screenshot of the therapist- and the patient-screen while perfoming a balance rehabilitation exercise. The playback speed has been lowered for this exercise as shown on the upper left corner.

Figure 3.21: Screenshot of the therapist- and the patient-screen while perfoming a rehabilitation exercise for the neck. The visual feedback indicates that the patient is raising the shoulder too much.

When the last exercise provided by the therapist and defined in the XML-file is reached and has been completed by the patient, the home-rehabilitation session is over and the patient has two possibilies: If the patient is satisfied with the scores reached, he or she can note the reached total score of the home-rehabilitation session and then close the application through the keyboard or the mouse. If the patient is not closing the application, the home-rehabilitation session is automatically started over again from the beginning, the scores are set back to zero and the patient can try to reach a higher score in the new rehabilitation session.

CHAPTER 4

Evaluation

In this chapter, the acquired results while evaluating the HomeRehab application are presented. Initially, the performance of the HomeRehab application in distinguishing between correct and incorrect rehabilitation exercise executions by the patients is examined and evaluated. For this performance test, Dr. Paternostro-Sluga provided six different typical rehabilitation exercises which represent the dataset for the evaluation and which are used as the reference exercises for the test. Based on this dataset, a home-rehabilitation session with these six precisely defined rehabilitation exercises, executed by three participants, has been performed and the results have been documented. The dataset and the outcomes of the performance evaluation for each exercise, as well as a summary of the classification results are presented in the following sections. Subsequently, the application is evaluated in regard to usability aspects. And finally, the observed limitations of the system are examined.

4.1 Dataset

In the course of a meeting at the *Sozialmedizinisches Zentrum Ost* in Vienna, six precisely defined rehabilitation exercises have been performed by a member of the staff of the *Institute of Physical Medicine and Rehabilitation* under supervision of Dr. Paternostro-Sluga. All six exercises have been executed for both the left and the right side of the body. The motions have been recorded and the resulting 12 exercise motion-files have been generated through the *NiViewer* application provided by *OpenNI* while using an *ASUS Xtion PRO LIVE* depth-sensor. Beside recording the correct exercise execution of the six different reference rehabilitation exercises, Dr. Paternostro-Sluga additionally provided the exact definition of each exercise, including exercise parameters and criteria for a correct exercise execution. She also pointed out the difference between correct and incorrect exercise executions for all six rehabilitation exercises, which is required for assessing the classification performance of the HomeRehab application. In the following, the six reference rehabilitation exercises representing the dataset are described in detail and the defined criteria for each exercise are shown.

4.1.1 Exercise 1 - Standing on one leg

The first reference exercise provided aims on improving the patients' balance ability. The initial position is an upright standing with the legs kept parallel to each other. The stomach and the pelvic floor have to be tensed by the patients at any time while performing the exercise. From the initial position, one leg is slightly lifted while the patients' pelvis has to be kept as straight as possible. This posture of standing on one leg has then to be held for five seconds. After the five seconds the leg is slowly lowered and brought back to the initial position again. After the defined required number of repeats for one leg has been performed by the patients, the exercise has to be repeated with the other leg. The criterion for a correct exercise execution is that the patients keep their hips straight and do not lean to the right or left side while standing on one leg. For that reason the only criterion defined in the XML-file containing the exercise parameters is the straightness-criterion 'hips'. Figure 4.1 shows the therapist while performing this rehabilitation exercise. The required exercise criterion is visualized through an overlay on the therapist's body. In Figure 4.1(a), the therapist is shown in the initial position of the exercise while (b) shows the therapist with a slightly lifted leg.

Figure 4.1: Reference rehabilitation exercise 1 (*Standing on one leg*) performed by the therapist and used for evaluating the HomeRehab application's exercise classification performance. The initial position of the exercise (a) as well as the posture to be held by the patients (b) are shown. While performing the exercise, the patients have to keep their hips as straight as possible.

4.1.2 Exercise 2 - One-legged squats

The second provided reference exercise is aimed on knee rehabilitation. The initial position is one leg stepped forward while the respective knee is stretched. The patients have to look straight, put their hands to the hips and keep their pelvis straight at any time of the exercise execution. From the initial position, the knee in the front is slowly bent while the patients' foot and the knee should remain vertically aligned at any time. The patients are not allowed to tilt away their knee to the left or right side while bending it. The posture with the bent knee has then to be held for 10 seconds while the knee has to be kept in position without tilting it away. After the 10 seconds, the knee is slowly stretched, again without tilting the knee to the side, until the initial position is reached again. The exercise is repeated with the same knee until the required number of repeats is reached. After that, the initial position is switched, meaning that the other foot is stepped forward, and the exercise is repeated with the patients' other leg. The criterion for a correct exercise execution of this rehabilitation exercise is that the patients keep the bent knee vertically in line with their foot and not tilt it to one side or the other at any time of the exercise execution. For that reason the only criterion defined in the XML-file containing the exercise parameters is the straightness-criterion 'right_leg' for the right body side and 'left_leg' for the left side respectively. In Figure 4.2 the therapist is shown while performing this rehabilitation exercise while the exercise criterion to be met is overlaid. The required initial position can be seen in Figure 4.2(a) while in (b) the therapist is shown while performing the squat.

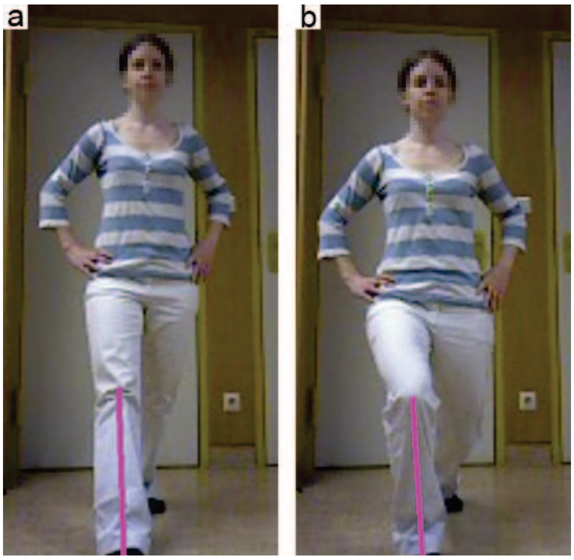

Figure 4.2: Reference rehabilitation exercise 2 (*One-legged squats*) performed by the therapist and used for evaluating the HomeRehab application's exercise classification performance. In (a) the initial position of the exercise is shown while (b) shows the squat and the posture to be held for 10 seconds. The patients are not allowed to tilt the knee to one side or the other at any time of the exercise execution.

4.1.3 Exercise 3 - Scapula exercise

The third reference exercise provided for the system evaluation is an exercise aimed on strengthening the patients' scapular muscles (shoulder blade muscles). The initial exercise position is an upright standing with the knees slightly bent. The patients have to look straight and tense their stomach and pelvic floor throughout the exercise. The upper arms are kept close to the body while the patients' elbows are bent to 90 degrees with the hand palms pointing up. The patients then have to slowly turn their arms from the front to the side for 60 degrees. This posture has then to be held for five seconds. Throughout the exercise, the upper arms have to be kept close to body and the shoulder blade muscles have to be pulled down and towards each other. The execution of this exercise by the therapist is shown in Figure 4.3 where (a) presents the initial position for the exercise and (b) the posture with the arms turned to the side while the scapular muscles are pulled down. After the five seconds, the arms are slowly turned back to the initial position. As visualized in the image, the criteria for a correct exercise execution are to keep the shoulders straight and to not put the upper arms away from the body at any time of the exercise. Both criteria can be checked by the application through defining the shoulders-straight and shoulder-angle criteria in the parameters XML-file.

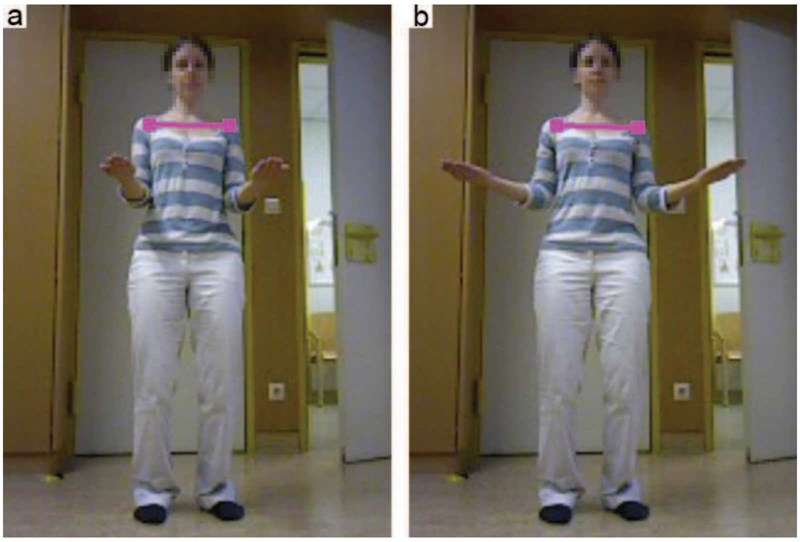

Figure 4.3: Reference rehabilitation exercise 3 (*Scapula exercise*) performed by the therapist and used for evaluating the HomeRehab application's exercise classification performance. The initial position (a) as well as the posture with the arms turned to the side (b) are shown. While performing the exercise, the patients are not allowed to raise a shoulder or to turn their upper arms away from their body.

4.1.4 Exercise 4 - Lift arm

The fourth provided reference rehabilitation exercise is aimed on shoulder and arm rehabilitation. The initial position is an upright standing with the knees slightly bent. The patients again have to tense their stomach and pelvic floor. The arm is constantly kept slightly bent while the patients' thumb points up. Figure 4.4(a) shows the therapist while being in the initial exercise position. The patients then have to slowly raise their arm above their head on the side of their body while keeping the shoulders as straight as possible. The shoulder blade muscles have to be pulled down and towards each other throughout the exercise execution. As soon as the arm is positioned above the head as shown in Figure 4.4(b), it is slowly lowered until the initial position is reached again. When the required number of repeats has been performed for this arm, the exercise has to be repeated with the other arm. In this exercise, three criteria are defined in the parameters XML-file. Beside the criterion to keep the shoulders straight, the elbow- and shoulder-angles have to remain within specific limits throughout the exercise execution. Again, the required exercise criteria are overlaid on the therapist's body in the shown image.

Figure 4.4: Reference rehabilitation exercise 4 (*Lift arm*) performed by the therapist and used for evaluating the HomeRehab application's exercise classification performance. In (a) the initial position of the exercise is shown while (b) shows the therapist while reaching the point where the arm is lowered again. The patients have to mimic the therapist's shoulder and elbow posture throughout the exercise execution and are not allowed to raise their shoulder while lifting the arm.

4.1.5 Exercise 5 - Cervical spine rotation

The fifth reference exercise provided is an exercise aimed on neck and cervical spine rehabilitation. The initial position, again, is an upright standing with the knees slightly bent and a tense stomach and pelvic floor. The patients have to look straight and relax their shoulders at any time of exercise execution. From the initial position, the head is slowly turned 45 to 60 degrees to one side. With the head turned, the patients have to slowly nod three times before the head is turned back to the initial position. The correct exercise execution is shown in Figure 4.5 with the initial exercise position presented in (a) and the rotated head before nodding shown in (b). After the required number of repeats for one side is reached, the exercise is repeated with turning the head to the other side. While turning the head to the side and while nodding with the turned head, the patients are not allowed to raise the shoulder. As this is the only criterion for a correct exercise execution in this rehabilitation exercise, the only criterion defined in the XML-file with the exercise parameters is the straightness-parameter 'shoulders'.

Figure 4.5: Reference rehabilitation exercise 5 (*Cervical spine rotation*) performed by the therapist and used for evaluating the HomeRehab application's exercise classification performance. The initial position (a) and the turned head of the therapist (b) are shown. After nodding three times while the head it is turned, it is slowly brought back to the initial position. While turning the head and while nodding with the turned head, the patients are not allowed to raise their shoulder.

4.1.6 Exercise 6 - Cervical spine tilt

The sixth and last reference exercise provided by Dr. Paternostro-Sluga is similar to exercise 5, with the difference that the head is not turned to the side but tilted instead. The initial position is the same as stated in exercise 5. From the initial position, the patients have to slowly tilt their head to one side while trying not to raise their shoulder on this side. The posture with the tilted head then has to be held for 10 seconds before the head is slowly brought back to the initial position again. After the required repeats for one side are reached, the exercise is repeated with tilting the head to the other side. The exercise criterion again is to keep the shoulders straight at any time of the exercise execution which again is defined in the parameters XML-file for this exercise. The *Cervical spine tilt* exercise execution by the therapist is shown in Figure 4.6 where the initial position (a) and the posture with the tilted head (b), as well as an overlay with the exercise criterion to be met, are presented.

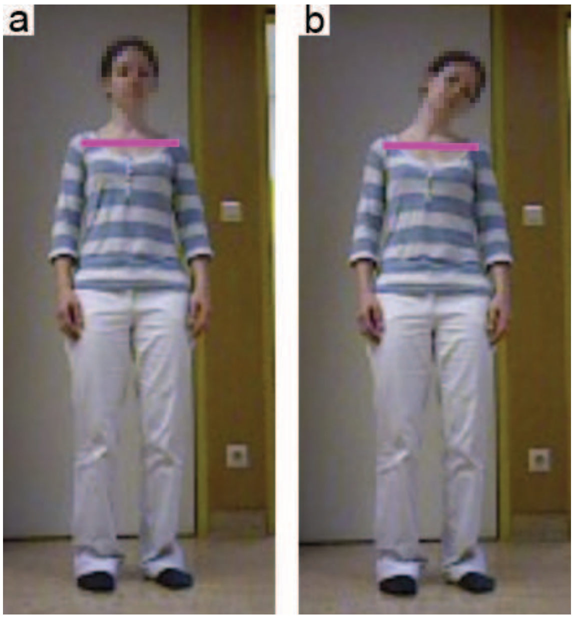

Figure 4.6: Reference rehabilitation exercise 6 (*Cervical spine tilt*) performed by the therapist and used for evaluating the HomeRehab application's exercise classification performance. The head is tilted from the initial position (a) to the side where it has to be held for 10 seconds (b). The patients are not allowed to raise the shoulder on the side where the head is tilted at any time of the exercise execution.

4.2 Classification Results

For testing the exercise classification performance of the HomeRehab application, the 12 motion-files (six exercises, each performed for the right and left side of the body) provided have been used as the reference for a correct exercise execution. Three subjects (two male, one female, ages 26 to 30, all healthy) participated in the test session representing the patients. The participants are referred to as *Subject 1*, *Subject 2* and *Subject 3* in the tables. In the executed home-rehabilitation session, every rehabilitation exercise has been performed for both sides of the body by every participant. The participants have been advised to perform each exercise 20 times for each side of the body, of which the first 10 repetitions have been executed correctly while the other 10 repetitions have been executed with specific mistakes, as defined in a meeting with Dr. Paternostro-Sluga. The aim of this approach has been to determine if the HomeRehab application is able to detect these mistakes in the exercise execution and thus classify the execution respectively. The tasks of every participant and thus the temporal sequence of the test session for every subject has been as follows:

- Perform the exercise on the right body side 10 times correctly
- Perform the exercise on the right body side 10 times incorrectly
- Perform the exercise on the left body side 10 times correctly
- Perform the exercise on the left body side 10 times incorrectly
- Proceed to the next exercise and repeat all steps until all six exercises are completed

Through this procedure, every participant performed a total of 120 exercise repetitions, leading to a total number of 360 repetitions of which 180 have been performed correctly while the other 180 have been performed incorrectly. The aim of the test is to evaluate the performance of the prototype application's ability to distinguish between correct and incorrect exercise executions. This ability can be checked through the scoring mechanism of the application, as an exercise execution classified as correct always leads to a maximum score of 12 points per repetition while exercise executions classified as incorrect always lead to lowered score values. The tolerance value for each exercise defined in the parameters XML-file has been set to 25% which has, due to the outcomes of first tests of the application, proven to be a value where healthy subjects are easily able to mimic the shown exercises correctly. The playback speed of the exercises has not been modified at any time of the test.

In the following, the results of the test rehabilitation session with the participants are presented separately for each exercise. Subsequently, the results of all exercises are summed up and the overall classification performance of the application is calculated and presented. And finally the calculated results are analyzed and discussed.

For every exercise repetition, the *True Positives* (TP), *False Negatives* (FN), *True Negatives* (TN) and *False Positives* (FP) the HomeRehab application generates have been examined and documented. True Positive values result if the subject performs the rehabilitation exercise correctly and the application classifies the exercise execution as correct. The better the classification performance of the application, the higher the number of True Positive values. False Negative values arise if a correctly performed exercise by the subject is classified as incorrect by the application. These values should be low as they indicate that the application is providing wrong feedback by notifying the subjects about mistakes although they are performing the exercises as they should. True Negative values indicate that the application successfully classifies incorrect exercise executions of the subjects as incorrect. These True Negative values are focused on most, as they represent the application's ability to notify the subjects about wrong exercise executions, which is the major aim of the HomeRehab project. A high True Negative rate is a result of a good classification performance of the application. And False Positive values result if the subject is making mistakes while performing the rehabilitation exercise and thus performs the exercise incorrectly but the application still classifies the exercise execution as correct. These values again should be avoided since they indicate that the application is not notifying the subjects about mistakes in their current exercise execution. Based on these results, the application's F_1-*score*, representing the classification performance of the application, is calculated and presented. The F_1-score is calculated through the harmonic mean of *Precision* and *Recall* as shown in Equation 4.1.

$$F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall} \tag{4.1}$$

The Precision is given by the number of True Positives divided by the sum of True Positives and False Positives. Recall is calculated through the number of True Positives divided by the sum of True Positives and False Negatives. In the following subsections, the classification results of the home-rehabilitation session are examined and the F_1-score is calculated for every exercise. Subsequently, the overall F_1-score is calculated and the results are discussed.

4.2.1 Exercise 1

The application's classification results in the test session for exercise 1 are shown in Table 4.1. The task has been to stand on one leg while the application checks if the subjects' hips are kept straight and through this if the subjects are leaning to the left or right side. As the results show, the application classified 96.7% of all correctly performed exercises as correct and 100% of wrong exercise executions as incorrect. The three False Negative values in this session resulted due to tracking errors of the depth-sensor, as parts of the subjects' skeleton has been lost for a few seconds at Subject 2 and Subject 3, leading to a wrong classification of three correctly executed exercise repetitions. On the other hand every wrong exercise execution has been successfully detected by the application. Based on these results, the application's F_1-score for exercise 1 is *0.97*.

Table 4.1: Results of the test session for exercise 1 (Standing on one leg)

	Raising right leg				**Raising left leg**			
	TP	**FN**	**TN**	**FP**	**TP**	**FN**	**TN**	**FP**
Subject 1	10	0	10	0	10	0	10	0
Subject 2	9	1	10	0	10	0	10	0
Subject 3	10	0	10	0	8	2	10	0
Sum	**29**	**1**	**30**	**0**	**28**	**2**	**30**	**0**
	(96.7%)	**(3.3%)**	**(100%)**	**(0%)**	**(93.3%)**	**(6.7%)**	**(100%)**	**(0%)**

4.2.2 Exercise 2

Table 4.2 shows the reached classification results in the test session for exercise 2. In this exercise, the subjects performed one-legged squats as described above. The results reveal a higher False Negative rate in comparison to the other exercises. A possible reason for this outcome is the system's reduced performance on tracking lower body parts. As a study presented in [15] reveals, the tracking accuracy of the depth-sensing camera as well as the motion estimation performance of the NiTE middleware is reduced when tracking lower body extremities, compared to other parts of the body. The position of the knee joint has been estimated incorrectly by the system at least once at every subject, resulting in a F_1-score of *0.87*. Even a subject's wrong exercise execution has been classified as performed correctly by the application due to the fact that the tracking of the knee position did not follow the actual position of the subject's knee while tilting it away which otherwise should have resulted in respective feedback and lowered score.

Table 4.2: Results of the test session for exercise 2 (One-legged squats)

	Bending right knee				**Bending left knee**			
	TP	**FN**	**TN**	**FP**	**TP**	**FN**	**TN**	**FP**
Subject 1	9	1	9	1	10	0	10	0
Subject 2	7	3	10	0	6	4	10	0
Subject 3	8	2	10	0	7	3	10	0
Sum	**24**	**6**	**29**	**1**	**23**	**7**	**30**	**0**
	(80%)	**(20%)**	**(96.7%)**	**(3.3%)**	**(76.7%)**	**(23.3%)**	**(100%)**	**(0%)**

4.2.3 Exercise 3

The classification results of the test session for exercise 3 can be found in Table 4.3. The exercise aimed on strengthening patients' scapular muscles proved to be very suitable for use with the HomeRehab application as a F_1-score of *1.0* has been reached in the test session. Every mistake of the subjects, including raising their shoulder or not keeping their upper arm close to the body, has been detected and reported by the application while every correct exercise execution has been classified as performed correctly.

4.2.4 Exercise 4

The classification results reached for exercise 4 are shown in Table 4.4. This exercise has been challenging for the participants as there are three criteria that have to be met while lifting the arms for the execution being classified as correct. The application analyzed if the subjects kept their shoulders straight and if the subjects' elbow- and shoulder-angles remained within the defined limits. The calculated F_1-score from the results of this exercise is *0.97*. The only incorrect classification arised when the subjects' estimated shoulder joint positions have been incorrect

Table 4.3: Results of the test session for exercise 3 (Scapula exercise)

	Tilting right arm				**Tilting left arm**			
	TP	**FN**	**TN**	**FP**	**TP**	**FN**	**TN**	**FP**
Subject 1	10	0	10	0	10	0	10	0
Subject 2	10	0	10	0	10	0	10	0
Subject 3	10	0	10	0	10	0	10	0
Sum	**30**	**0**	**30**	**0**	**30**	**0**	**30**	**0**
	(100%)	**(0%)**	**(100%)**	**(0%)**	**(100%)**	**(0%)**	**(100%)**	**(0%)**

for a few frames, leading to the consequence that the criterion to keep the shoulders straight is not met anymore, although the subjects did not raise their shoulders. Short tracking errors while analyzing the elbow- and shoulder-angles did not influence the result due to the fact that these criteria are still met if the calculated angles are beyond the defined limits for less than one second, as described in Chapter 3. Again, as no False Positive values have been generated, every wrong exercise execution of the subjects has been detected and classified as performed incorrectly by the application.

Table 4.4: Results of the test session for exercise 4 (Lift arm)

	Lifting right arm				**Lifting left arm**			
	TP	**FN**	**TN**	**FP**	**TP**	**FN**	**TN**	**FP**
Subject 1	10	0	10	0	9	1	10	0
Subject 2	10	0	10	0	10	0	10	0
Subject 3	9	1	10	0	8	2	10	0
Sum	**29**	**1**	**30**	**0**	**27**	**3**	**30**	**0**
	(96.7%)	**(3.3%)**	**(100%)**	**(0%)**	**(90%)**	**(10%)**	**(100%)**	**(0%)**

4.2.5 Exercise 5

Table 4.5 presents the results of the application's exercise classification performance while performing exercise 5 in the test session. The only criterion the application has to check in the cervical spine rotation exercise is that the subjects are keeping their shoulders straight at any time while performing the exercise. The classification results show that every exercise repetition of the subjects has been classified correctly, meaning that every wrong shoulder posture of the subjects has been detected. Beside that, the tracking and estimating of the subjects' shoulder positions has never been interfered throughout the test session on this exercise. Thus, the resulting F_1-score of the application for the cervical spine rotation exercise in the test session has reached *1.0*.

Table 4.5: Results of the test session for exercise 5 (Cervical spine rotation)

	Rotating head to the right				**Rotating head to the left**			
	TP	**FN**	**TN**	**FP**	**TP**	**FN**	**TN**	**FP**
Subject 1	10	0	10	0	10	0	10	0
Subject 2	10	0	10	0	10	0	10	0
Subject 3	10	0	10	0	10	0	10	0
Sum	**30**	**0**	**30**	**0**	**30**	**0**	**30**	**0**
	(100%)	**(0%)**	**(100%)**	**(0%)**	**(100%)**	**(0%)**	**(100%)**	**(0%)**

4.2.6 Exercise 6

The classification results of the test session for exercise 6 are shown in Table 4.6. For the HomeRehab application, the classification task in this exercise is the same as in exercise 5 due to the fact that the only exercise criterion for a correct exercise execution is that the subjects do not raise their shoulders while tilting their head. However, for the subjects a tilt of the head without raising the shoulders is harder to perform than a rotation. The only False Negative value, appearing at Subject 2 can again be explained through a wrong shoulder position estimation of the system, which lasted about one second and lead the application to classify a correctly performed exercise repetition as incorrect. The F_1-score for this exercise is *0.99*.

Table 4.6: Results of the test session for exercise 6 (Cervical spine tilt)

	Tilting head to the right				**Tilting head to the left**			
	TP	**FN**	**TN**	**FP**	**TP**	**FN**	**TN**	**FP**
Subject 1	10	0	10	0	10	0	10	0
Subject 2	9	1	10	0	10	0	10	0
Subject 3	10	0	10	0	10	0	10	0
Sum	**29**	**1**	**30**	**0**	**30**	**0**	**30**	**0**
	(96.7%)	**(3.3%)**	**(100%)**	**(0%)**	**(100%)**	**(0%)**	**(100%)**	**(0%)**

4.3 Summary of the Classification Results

This section summarizes the classification results collected through the performed home rehabilitation session including six rehabilitation exercises performed by three participants for both sides of the body. The application's results for the performed exercises are summarized in Table 4.7. The values presented in this table for each exercise are averaged for the exercise execution on the right and the left body side.

In this rehabilitation session, an overall True Positives rate of 94.2% has been reached by the application, meaning that 94.2% of the 180 correctly performed exercise repetitions have been classified as executed correctly. From the 180 incorrect exercise executions performed by the participants, a total of 99.7% have been classified as incorrect by the application, as can be seen at the True Negatives rate in the table. These results lead to an overall F_1-score of the HomeRehab application of **0.967** for the six reference rehabilitation exercises provided by Dr. Paternostro-Sluga.

This result is very promising and confirms the application's possibility of successfully support patients in their home-rehabilitation sessions. The evaluated rehabilitation session revealed that the main reasons for wrong classifications are interferences in the tracking of the depth-sensor which are causing short aberrations of the estimated skeletal joint positions from the real positions of the subjects' skeletal joints. Through a reduction of these interferences, the correct classification rate of the application could even be higher.

Table 4.7: Summarized results for each performed exercise in the home-rehabilitation session

	TP (%)	**FN (%)**	**TN (%)**	**FP (%)**
Exercise 1	95 %	5 %	100 %	0 %
Exercise 2	78.4 %	21.6 %	98.4 %	1.6 %
Exercise 3	100 %	0 %	100 %	0 %
Exercise 4	93.4 %	6.6 %	100 %	0 %
Exercise 5	100 %	0 %	100 %	0 %
Exercise 6	98.4 %	1.6 %	100 %	0 %
Average	**94.2 %**	**5.8 %**	**99.7 %**	**0.3 %**

4.4 Usability Evaluation

Beside the application's performance in distinguishing between correct and incorrect exercise executions, the application has been evaluated in regard to usability as well. On the one hand, the usability of the HomeRehab application is evaluated through the feedback of Dr. Paternostro-Sluga regarding the interaction with the application, acquired in the course of the meetings at the *Sozialmedizinisches Zentrum Ost*. On the other hand, the participants of the application's classification performance test have been questioned about interaction and usability aspects after completing the test rehabilitation session. Based on these informations, the following statements about the HomeRehab application's usability can be given:

The overall impression of the application interface and the interactive character of the application have been outstanding. The textual guidance the application provides to the patients enables a very intuitive use and easy comprehensibility. Also the generated visual feedback regarding the exercise execution has been understood without further explanation. However, the actual interaction with the application while performing the rehabilitation exercise requires a short familiarization phase, as the patients are required to alternatingly check the therapist-screen for the exercise progress and the patient-screen for the visual feedback.

The calculated and visualized scores has proven to be useful to motivate the users for a more accurate exercise execution and for constantly trying to raise their personal highscores.

Due to the fact that the patient has to keep a specific minimum distance to the depth-sensing camera for proper motion tracking, and to the fact that ideally no other objects are occluding the patients' body in the tracked scene, operating the application on a laptop is impracticable because it has to be placed outside of the sensor's view which can hinder a proper usage of the application because of the reduced screen size. The application is intended to be operated with a large screen, or ideally a TV screen, and with a depth-sensor placed underneath or above this screen.

In addition, Dr. Paternostro-Sluga suggested a second operational area in which the HomeRehab application could prove to be useful. The application could additionally be used for motivating elderly people to physical activity. According to the World Health Organization (WHO), continuous physical exercise in old age leads to fewer cases of stroke, reduced rates of coronary heart disease, reduced mortality rates and an overall health enhancement [17]. For using the HomeRehab application for this purpose it is sufficient to simply provide motion-files containing movements that elderly people are able to perform easily and set the tolerance value in the parameters file to 100%. Through this approach the motions of the elderly people are tracked by the depth-sensor and visualized on the screen but not analyzed and checked for correctness by the system.

4.5 Limitations of the HomeRehab-System

The full functionality of the HomeRehab application is dependent on accurate tracking of the patients' motions and a reliable reconstruction of the patients' skeletal joint positions. Due to the fact that the OpenNI framework in combination with the NiTE middleware is not able to reconstruct the skeleton of a tracked human body if the lower half of the body is occluded, the HomeRehab application is not suitable for rehabilitation exercises where the patients are in a sitting position. Therefore, also home-rehabilitation sessions with patients in wheelchairs can not be supported by this application. Rehabilitation exercises to be performed in lying positions are not suitable either, as a proper skeleton tracking is not possible in these cases.

Another limitation of the usage of the application are rehabilitation exercises based on motions that are not detectable by the depth-sensing camera. For instance, exercises where specific muscles have to be tensed without actually moving the respective body parts.

Beside that, since the depth-sensor in this project is based on infrared light, direct sunlight hitting the sensor or the tracked scene can interfere the sensor and impede proper motion tracking. Although the system is designed for indoor-use only, even sunlight through a window can be problematic as experienced while testing the HomeRehab application.

CHAPTER 5

Discussion and Conclusion

In the course of this master's thesis, a prototype application for supporting home-based physical rehabilitation has been developed, implemented and evaluated. The application is based on a low-cost depth-sensing camera for tracking the patients' motions while performing rehabilitation exercises at home and is providing guidance and visual feedback regarding the correctness of the exercise execution. What distinguishes the HomeRehab application from other home-based computer-assisted rehabilitation projects is the possibility of using arbitrarily defined rehabilitation exercises, specifically adjusted according to the current conditions of the patients. Through the association with a medical professional in physical medicine and rehabilitation during project development and evaluation, the application is considering the needs of physically disabled persons.

The evaluation of the project revealed promising results in the performance of distinguishing between correct and incorrect exercise executions and in usability aspects. The major task of the application is to prevent patients from wrong rehabilitation exercise executions at their home, since no therapist is supervising them and notifying them about mistakes when performing the exercises at home. As the performed evaluation showed, 99.7% of wrong exercise executions have been detected and classified as incorrect, leading to respective notifications via the visual feedback. This promising result proves the suitability of the HomeRehab application for the use of preventing patients from wrong exercise executions if the specifications of the rehabilitation exercises are precisely defined by a therapist. As the overall correct classification rate of 96.95% is mainly reduced by correct exercise executions classified by the application as performed incorrectly, due to tracking errors of the depth-sensor, compensating these errors would lead to an almost perfect classification performance.

In regard to usability and acceptance, the results are very promising as well. Due to the clear and simple interface and the textual guidance, the application proved to be very intuitive to use. Also the expected beneficial effects of the scoring mechanism for increased application usage motivation have occurred. Additionally, through the different adjustable exercise parameters, the rehabilitation exercises can continuously be tailored to the patients' current condition. By lowering the tolerance value of the exercises according to the patients' rehabilitation progress

for example, the exercises can continuously be kept challenging to prevent the patients from a descending motivation for performing the required rehabilitation exercises.
A next step would be a clinical trial with patients suffering from different motor disabilities using the HomeRehab application. Beside that, several options for possible future improvements are conceivable: A useful extension to the application would be an user interface providing the possibility to navigate through different options like skipping an exercise, starting over from the beginning or choosing which exercise is to be loaded next. This interface could be controlled through hand gestures or, due to the audio recording possibility of the depth-sensing camera, through spoken instructions via speech recognition. Another possible extension would be a self-adaptive approach where the playback speed and the tolerance regarding the exercise execution accuracy are adjusted automatically by interpreting the patients' current motions in regard to their current condition. Beside that, an additional therapist application could be implemented, providing the therapists with the possibility to communicate with the applications of their patients to comfortably provide new exercises and adjust exercise parameters remotely, and additionally monitor the patients' rehabilitation progress.

Bibliography

[1] Mohamed Abdur Rahman, Ahmad M. Qamar, Mohamed A. Ahmed, M. Ataur Rahman, and Saleh Basalamah. Multimedia interactive therapy environment for children having physical disabilities. In *Proceedings of the 3rd ACM International Conference on Multimedia Retrieval*, pages 313–314. ACM, 2013.

[2] Dimitrios S. Alexiadis, Philip Kelly, Petros Daras, Noel E. O'Connor, Tamy Boubekeur, and Maher Ben Moussa. Evaluating a dancer's performance using kinect-based skeleton tracking. In *Proceedings of the 19th ACM International Conference on Multimedia*, pages 659–662. ACM, 2011.

[3] Simon Baker, Takeo Kanade, and Kong-Man Cheung. Shape-from-silhouette across time part 1: Theory and algorithms. *International Journal of Computer Vision*, 62(3):221–247, 2005.

[4] Antonio Bo, Mitsuhiro Hayashibe, and Philippe Poignet. Joint angle estimation in rehabilitation with inertial sensors and its integration with kinect. In *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC*, pages 3479–3483, August 2011.

[5] Hans-Peter Bruckner, Rochus Nowosielski, Henning Kluge, and Holger Blume. Mobile and wireless inertial sensor platform for motion capturing in stroke rehabilitation sessions. In *5th IEEE International Workshop on Advances in Sensors and Interfaces 2013 (IWASI)*, pages 14–19, June 2013.

[6] Grigore Burdea, Viorel Popescu, Vincent Hentz, and Kerri Colbert. Virtual reality-based orthopedic telerehabilitation. *IEEE Transactions on Rehabilitation Engineering*, 8(3):430–432, September 2000.

[7] James William Burke, Michael McNeill, Darryl Charles, Philip J. Morrow, Jacqui Crosbie, and Suzanne McDonough. Serious games for upper limb rehabilitation following stroke. In *Conference in Games and Virtual Worlds for Serious Applications, VS-GAMES'09*, pages 103–110. IEEE, 2009.

[8] Fabrice Caillette and Toby Howard. Real-time markerless human body tracking using colored voxels and 3d blobs. In *Third IEEE and ACM International Symposium on Mixed and Augmented Reality, ISMAR 2004*, pages 266–267. IEEE, 2004.

[9] Chien-Yen Chang, Belinda Lange, Mi Zhang, Sebastian Koenig, Phil Requejo, Noom Somboon, Alexander A. Sawchuk, and Albert A. Rizzo. Towards pervasive physical rehabilitation using microsoft kinect. In *2012 6th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, pages 159–162. IEEE, 2012.

[10] Arie Croitoru, Peggy Agouris, and Anthony Stefanidis. 3d trajectory matching by pose normalization. In *Proceedings of the 13th Annual ACM International Workshop on Geographic Information Systems*, pages 153–162. ACM, 2005.

[11] Alana Da Gama, Thiago Chaves, Lucas Figueiredo, and Veronica Teichrieb. Guidance and movement correction based on therapeutics movements for motor rehabilitation support systems. In *Proceedings of the 14th Symposium on Virtual and Augmented Reality 2012*, pages 191–200. IEEE Computer Society, 2012.

[12] Alana Da Gama, Thiago Chaves, Lucas Figueiredo, and Veronica Teichrieb. Improving motor rehabilitation process through a natural interaction based system using kinect sensor. In *2012 IEEE Symposium on 3D User Interfaces (3DUI)*, pages 145–146, March 2012.

[13] Ruth Dickstein, Shraga Hocherman, Thomas Pillar, and Rachel Shaham. Stroke rehabilitation - three exercise therapy approaches. *Physical Therapy*, 66(8):1233–1238, 1986.

[14] Pamela Duncan, Lorie Richards, Dennis Wallace, Joni Stoker-Yates, Patricia Pohl, Carl Luchies, Abna Ogle, and Stephanie Studenski. A randomized, controlled pilot study of a home-based exercise program for individuals with mild and moderate stroke. *Stroke*, 29(10):2055–2060, 1998.

[15] Adso Fernandez-Baena, Antonio Susin, and Xavier Lligadas. Biomechanical validation of upper-body and lower-body joint movements of kinect motion capture data for rehabilitation treatments. In *2012 4th International Conference on Intelligent Networking and Collaborative Systems (INCoS)*, pages 656–661. IEEE, 2012.

[16] Peter Fikar, Christian Schoenauer, and Hannes Kaufmann. The sorcerer's apprentice: A serious game aiding rehabilitation in the context of subacromial impingement syndrome. In *Proceedings of the 7th International Conference on Pervasive Computing Technologies for Healthcare*, pages 327–330. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2013.

[17] Samyukta Ganesan and Lisa Anthony. Using the kinect to encourage older adults to exercise: A prototype. In *CHI'12 Extended Abstracts on Human Factors in Computing Systems*, pages 2297–2302. ACM, 2012.

[18] Juan Enrique Garrido, Irma Marset, Victor M.R. Penichet, and Maria D. Lozano. Balance disorder rehabilitation through movement interaction. In *2013 7th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, pages 319–322, May 2013.

[19] Marientina Gotsis, Fotos Frangoudes, Vangelis Lympouridis, Somboon Maneekobkunwong, David Turpin, and Maryalice Jordan-Marsh. Skyfarer: A mixed reality shoulder exercise game. In *ACM SIGGRAPH 2013 Studio Talks*, pages 2:1–2:1. ACM, 2013.

[20] Thomas Helten. *Processing and tracking human motions using optical, inertial, and depth sensors*. PhD thesis, Universitaet des Saarlandes, Germany, 2013.

[21] Jun-Da Huang. Kinerehab: A kinect-based system for physical rehabilitation - a pilot study for young adults with motor disabilities. In *The Proceedings of the 13th International ACM SIGACCESS Conference on Computers and Accessibility*, pages 319–320. ACM, 2011.

[22] Zhifeng Huang, Ayanori Nagata, Masako Kanai-Pak, Jukai Maeda, Yasuko Kitajima, Mitsuhiro Nakamura, Kyouko Aida, Noriaki Kuwahara, Taiki Ogata, and Jun Ota. Development of a nursing self-training system for transferring patient from bed to wheelchair. In *2012 Proceedings of SICE Annual Conference (SICE)*, pages 246–254, August 2012.

[23] Feifei Huo, Emile Hendriks, Pavel Paclik, and Augustinus H.J. Oomes. Markerless human motion capture and pose recognition. In *10th Workshop on Image Analysis for Multimedia Interactive Services, WIAMIS'09*, pages 13–16. IEEE, 2009.

[24] Mohamed E. Hussein, Marwan Torki, Mohammad A. Gowayyed, and Motaz El-Saban. Human action recognition using a temporal hierarchy of covariance descriptors on 3d joint locations. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, pages 2466–2472. AAAI Press, 2013.

[25] Helmut Ivansits. Gesundheitspolitische dimensionen und ziele der beruflichen rehabilitation in oesterreich. *WISO - Wirtschafts- und Sozialpolitische Zeitschrift*, 4:15–28, April 2012.

[26] Carolin Jakob, Patrick Kugler, Felix Hebenstreit, Samuel Reinfelder, Ulf Jensen, Dominik Schuldhaus, Matthias Lochmann, and Bjoern M. Eskofier. Estimation of the knee flexion-extension angle during dynamic sport motions using body-worn inertial sensors. In *Proceedings of the 8th International Conference on Body Area Networks*, pages 289–295. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2013.

[27] Kate Jolly, Rod S. Taylor, Gregory Y.H. Lip, and Andrew Stevens. Home-based cardiac rehabilitation compared with centre-based rehabilitation and usual care: A systematic review and meta-analysis. *International Journal of Cardiology*, 111(3):343–351, 2006.

[28] Han Jungong, Shao Ling, Xu Dong, and J. Shotton. Enhanced computer vision with microsoft kinect sensor: A review. *IEEE Transactions on Cybernetics*, 43(5):1318–1334, October 2013.

[29] Kourosh Khoshelham. Accuracy analysis of kinect depth data. In *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences ISPRS laser scanning workshop*, volume 38, pages 133–138, August 2011.

[30] Jihyoung Kim, Sungwon Yang, and Mario Gerla. Stroketrack: Wireless inertial motion tracking of human arms for stroke telerehabilitation. In *Proceedings of the First ACM Workshop on Mobile Systems, Applications, and Services for Healthcare*, pages 4:1–4:6. ACM, 2011.

[31] Belinda Lange, Chien-Yen Chang, Evan Suma, Bradley Newman, Albert Skip Rizzo, and Mark Bolas. Development and evaluation of low cost game-based balance rehabilitation tool using the microsoft kinect sensor. In *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC*, pages 1831–1834, August 2011.

[32] Antoine Lejeune, Sebastien Pierard, Mark Van Droogenbroeck, and Jacques Verly. A new jump edge detection method for 3d cameras. In *International Conference on 3D Imaging (IC3D)*, December 2011.

[33] Mark A. Livingston, Jay Sebastian, Zhuming Ai, and Jonathan W. Decker. Performance measurements for the microsoft kinect skeleton. In *2012 IEEE Virtual Reality Short Papers and Posters (VRW)*, pages 119–120. IEEE, 2012.

[34] Jose-Antonio Lozano-Quilis, Herme Gil-Gomez, Jose-Antonio Gil-Gomez, Sergio Albiol-Perez, Garcia Palacios, Habib M. Fardoum, and Abdulfattah S. Mashat. Virtual reality system for multiple sclerosis rehabilitation using kinect. In *Proceedings of the 7th International Conference on Pervasive Computing Technologies for Healthcare*, pages 366–369. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2013.

[35] Dario Maggiorini, Laura Ripamonti, and Andrea Scambia. Videogame technology to support seniors. In *Proceedings of the 5th International ICST Conference on Simulation Tools and Techniques*, pages 270–277. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2012.

[36] Zoe Marquardt, Joao Beira, Natalia Em, Isabel Paiva, and Sebastian Kox. Super mirror: A kinect interface for ballet dancers. In *CHI '12 Extended Abstracts on Human Factors in Computing Systems*, pages 1619–1624. ACM, 2012.

[37] Javier Martin-Moreno, Daniel Ruiz-Fernandez, Antonio Soriano-Paya, and Vicente Jesus Berenguer-Miralles. Monitoring 3d movements for the rehabilitation of joints in physiotherapy. In *30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS 2008*, pages 4836–4839, August 2008.

[38] Stepan Obdrzalek, Gregorij Kurillo, Ferda Ofli, Ruzena Bajcsy, Edmund Seto, Holly Jimison, and Michael Pavel. Accuracy and robustness of kinect pose estimation in the context of coaching of elderly population. In *2012 annual international conference of the IEEE Engineering in medicine and biology society (EMBC)*, pages 1188–1193. IEEE, 2012.

[39] Gerald T. O'Connor, Julie E. Buring, Salim Yusuf, Samuel Z. Goldhaber, Elaine M. Olmstead, R.S. Paffenbarger, and Charles H. Hennekens. An overview of randomized trials of rehabilitation with exercise after myocardial infarction. *Circulation*, 80(2):234–244, 1989.

[40] Michele Pirovano, Renato Mainetti, Gabriel Baud-Bovy, Pier Luca Lanzi, and Alberto N. Borghese. Self-adaptive games for rehabilitation at home. In *2012 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 179–186, September 2012.

[41] Christian Plagemann, Varun Ganapathi, Daphne Koller, and Sebastian Thrun. Real-time identification and localization of body parts from depth images. In *2010 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3108–3113. IEEE, 2010.

[42] Roland PS van Peppen. *Towards evidence-based physiotherapy for patients with stroke*. PhD thesis, Ultrecht University and Hogeschool Utrecht, The Netherlands, 2008.

[43] Michalis Raptis, Darko Kirovski, and Hugues Hoppe. Real-time classification of dance gestures from skeleton animation. In *Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 147–156. ACM, 2011.

[44] Natalia Diaz Rodriguez, Robin Wikstroem, Johan Lilius, Manuel Pegalajar Cuellar, and Miguel Delgado Calvo Flores. Understanding movement and interaction: An ontology for kinect-based 3d depth sensors. In *Ubiquitous Computing and Ambient Intelligence. Context-Awareness and Context-Driven Interaction*, pages 254–261, 2013.

[45] Anil K. Roy, Yash Soni, and Sonali Dubey. Enhancing effectiveness of motor rehabilitation using kinect motion sensing technology. In *2013 IEEE Global Humanitarian Technology Conference: South Asia Satellite (GHTC-SAS)*, pages 298–304, August 2013.

[46] Jamie Shotton, Toby Sharp, Alex Kipman, Andrew Fitzgibbon, Mark Finocchio, Andrew Blake, Mat Cook, and Richard Moore. Real-time human pose recognition in parts from single depth images. *Communications of the ACM*, 56(1):116–124, 2013.

[47] Stephen Simmons, Rachel McCrindle, Malcolm Sperrin, and Andy Smith. Prescription software for recovery and rehabilitation using microsoft kinect. In *2013 7th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*, pages 323–326, May 2013.

[48] Heidi Sveistrup. Motor rehabilitation using virtual reality. *Journal of NeuroEngineering and Rehabilitation*, 1(1), 2004.

[49] Norman Villaroman, Dale Rowe, and Bret Swan. Teaching natural user interaction using openni and the microsoft kinect sensor. In *Proceedings of the 2011 Conference on Information Technology Education*, pages 227–232. ACM, 2011.

[50] Lu Xia, Chia-Chih Chen, and Jake Aggarwal. Human detection using depth information by kinect. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2011 IEEE Computer Society Conference on*, pages 15–22. IEEE, 2011.

[51] Huiyu Zhou and Huosheng Hu. Inertial motion tracking of human arm movements in stroke rehabilitation. In *2005 IEEE International Conference on Mechatronics and Automation*, volume 3, pages 1306–1311. IEEE, 2005.

[52] Huiyu Zhou and Huosheng Hu. Human motion tracking for rehabilitation: A survey. *Biomedical Signal Processing and Control*, 3(1):1–18, 2008.