Autonomous Robots, 4 (2017) 3, 679-698, Springer 2017
The final publication is available at link.springer.com

ROREAS: robot coach for walking and orientation training in clinical post-stroke rehabilitation - prototype implementation and evaluation in field trials

H.-M. Gross · A. Scheidig · K. Debes · E. Einhorn · M. Eisenbach · St. Mueller · Th. Schmiedel · T. Trinh · Ch. Weinrich · T. Wengefeld · A. Bley · Ch. Martin

Received: 27 March 2015 / Accepted: 7 February 2016 / Published online: 24 February 2016

Abstract This paper describes the objectives and the state of implementation of the ROREAS project which aims at developing a socially assistive robot coach for walking and orientation training of stroke patients in the clinical rehabilitation. The robot coach is to autonomously accompany the patients during their exercises practicing their mobility skills. This requires strongly user-centered, polite and attentive social navigation and interaction abilities that can motivate the patients to start, continue, and regularly repeat their self-training. The paper gives an overview of the training scenario and describes the constraints and requirements arising from the scenario and the operational environment. Moreover, it presents the mobile robot ROREAS and gives an overview of the robot's system architecture and the required human- and situation-aware navigation and interaction skills. Finally, it describes our three-stage approach in conducting function and user tests in the clinical environment: pre-tests with technical staff, followed by function tests with clinical staff and user trials with volunteers from the group of stroke patients, and presents the results of these tests conducted so far.

Keywords Robotic rehabilitation assistant · Walking coach · Socially assistive robotics · Post-stroke rehabilitation

H.-M. Gross, A. Scheidig, K. Debes, E. Einhorn, M. Eisenbach, St. Mueller, Th. Schmiedel, T. Trinh, Ch. Weinrich and T. Wengefeld are with the Neuroinformatics and Cognitive Robotics Lab, Technische Universitaet Ilmenau, Germany.
E-mail: Horst-Michael.Gross@tu-ilmenau.de

A. Bley and Ch. Martin are with the MetraLabs Robotics GmbH Ilmenau, Germany.

1 Introduction

As motor and cognitive learning are not passive processes, patients recovering from a stroke must play an active role in the rehabilitation process if improvement is to occur (Andrade et al. 2014). Against this background, a new trend in rehabilitation care is promising vast medical as well as economic potential - the so-called self-training of the patients. This finding was the context and the motivation for the research project ROREAS (Gross et al. 2014) running from mid 2013 till the beginning of 2016, which aims at developing a robotic rehabilitation assistant for walking self-training of stroke patients in late stages of the clinical post-stroke rehabilitation. The robotic rehab assistant is to accompany patients who already got the permission to walk on their own without professional assistance during their walking exercises, practicing both mobility and spatial orientation skills. It shall also address the patients' insecurity and anxiety ("Am I able to do that", "Will I find my way back?") which are possible reasons for not performing or neglecting self-training. The assistant is also supposed to monitor and document the exercises and store clinical records for accounting and clearing with insurance funds, thus combining improved training capabilities for patients and organizational efficiency for the rehabilitation clinic.

The project requires consistent integration of intuitive assistive functions allowing customized individual exercise plans, advanced human-robot-interaction (HRI) skills, and robust and polite autonomous navigation in populated public environments. Beside the user-centered development and implementation of the robotic training assistant, comprehensive user tests with volunteers from the group of stroke patients and a detailed analysis of the results shall quantify its medical

Fig. 1 Robotic walking coach "ROREAS" during a walking tour in our test site, the "m&i Fachklinik" rehabilitation center in Bad Liebenstein (Germany)

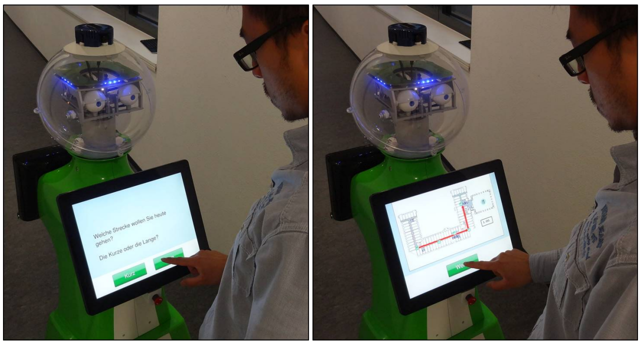

Fig. 2 GUI-based interaction between patient and walking coach "ROREAS". (left) intro-question regarding the desired length of a walking tour. (right) after training: floor plan with the route walked

effectiveness and reveal factors promoting or impeding the acceptance of its application.

Building on preceding own projects in socially assistive robotics for public operation (Gross et al. 2009) and domestic use (Gross et al. 2012), (Gross et al. 2015), the aim of the ROREAS project is (i) to complete the spectrum of robotic functionalities and services that are required for a robotic walking coach, and (ii) to evaluate the usability, the usefulness, and the added value of the rehab assistant for the patients during their clinical post-stroke rehabilitation.

The remainder of the article is organized as follows: Sect. 2 first gives an overview of the training scenario and describes the constraints and requirements arising from this specific rehabilitation scenario, while Sect. 3 discusses related work in the field of mobile rehabilitation robotics with a focus on walking training. Based on this, Sect. 4 presents the ROREAS prototype, an application-tailored mobile robot developed within the ROREAS project to meet the requirements to a personal training robot. Then Sect. 5 gives a brief overview of the robot's functional system architecture, while Sect. 6 introduces essential HRI and navigation skills required for a robot coach that can operate autonomously in such a challenging real-world environment like a rehab center. Using these functionalities, Sect. 7 introduces our three-stage approach in conducting the function and user tests in the clinical environment, presents first encouraging results of these tests conducted under clinical everyday conditions with volunteers from the group of stroke patients, and gives an outlook on upcoming user studies. Finally, Sect. 8 summarizes our main contributions.

2 Robotic walking coach in the rehab scenario

In this section, we outline a typical training session with the robotic walking coach, describe the specifics and challenges of the clinical setting, and define the resulting technical requirements for the robotic walking coach.

2.1 Typical training session

During the specification phase of the project, it was defined that only physically impaired stroke patients who already got the permission by their doctor in charge to walk on their own without professional assistance (usually using walking aids) are to be involved in this robot-assisted *Walking training*. The initially intended robot-assisted *Orientation training* of cognitively impaired stroke patients who require a dedicated training of their cognitive skills, e.g. by spatial exploration of the whole building or elevator usage, was completely postponed to a follow-up project due to the high complexity and the medical and ethical challenges and restrictions of this type of self-training. The duration of a walking training session was limited to no more than 20-30 minutes, depending of the physical conditions of the respective patient. However, it was allowed that multiple sessions can be scheduled on the same day to intensify the training.

In the following, a sketch of a typical walking training session is given as it is supposed to be realized till the final user tests at the end of the project. The training session is initiated either by the robot by sending a text message to the phone in the patient's room or optionally by the patient, who can call the robot by telephone. The robot then autonomously drives to the patient's room and takes a non-blocking waiting position at the door. It observes the corridor for a person emerging from the respective room door and then starts a verbal greeting, for instance: "Good morning, I am your personal walking coach. Please touch my screen." Then the patient logs in via touching a start button on the screen as asked by the robot. After that, the re-identification module on the robot learns the current appearance of the patient based on the present clothing, which is required for keeping track of the right user during a training session in a crowded environment.

Based on the training progress in preceding training sessions, the patient is offered a suggestion for the upcoming walk by means of speech and GUI output, for example: “On my screen you can see your last walk. You did the way to the first resting point twice. I think you can make it to the second resting point today.” Then the patient has to select a way (see Fig. 2, left) to walk or can cancel the session, if s/he feels too weak. After that, the robot gives instructions on the way and follows the patient: “Let’s go to the first target point. I will follow you. Touch my screen, if you need my help.”

If the patient takes the wrong way at the beginning or during a tour, the robot, which is following in a distance of two to three meters, detects this, points the patient on this issue, and waits for the patient to go the right way. If the patient does not react accordingly, the training session is canceled and the nursing staff gets informed by means of a text message.

On the way during the walking session, the robot points out possibilities for having a rest (the resting places ‘R’ in Fig. 3) and also remarks orientation features (e.g. pictures on the wall, plants, etc.) which are helpful for finding the way back on longer tours. Thus, the patient can either go on or take a seat to revive. If the robot detects, that the user has sat down too often, it suggests to finish the training and offers going back to the patient’s room or calling for the nurse. Also midway of the planned walk, the robot mentions that and suggest to return.

Depending on the interval of breaks and the distance gone, the robot reacts by means of motivating and encouraging speech and GUI outputs, such as: “Compared to our last session, you went a much longer way before you needed a break. This is a great improvement. Keep on going!” At the end, back in front of the patient’s room, the robot offers to continue the training and walk a bit further, if there is time left, or it summarizes the training (see Fig. 2, right), reminds the next scheduled appointment, and says good bye, for example: “We are back at your room. Today you made 50 meters. That is 10 meters more than last time. I will be here for our next training tomorrow at 5 pm, if you like. Good bye!”

“Guiding” and “Following” patients are two different modes of accompanying the patients by the robot. In both modes, the patients are physically active as they are walking - with or without walking aids - along the corridors of the rehab center. However, considering the cognitive difficulty, the “Following” mode is

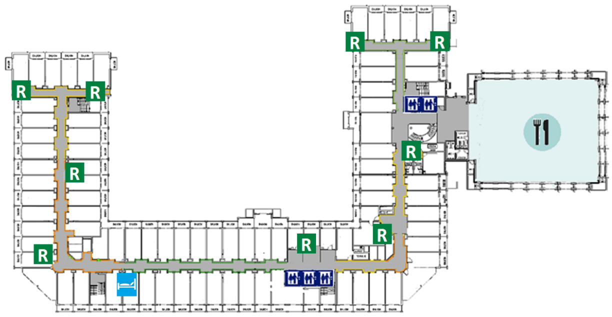

Fig. 3 Plan of one floor of the eight-floors rehabilitation center (m&i Fachklinik) in Bad Liebenstein (Germany) used as test site in the ROREAS project. The length of the corridor is about 170 meters. R marks resting places for the patients during their walking training along the corridor

more demanding for the patients than “Guiding”, as the patients have to take the initiative and walk in front of the robot which is only observing them and their walking behavior. Thus, the patients need to orientate themselves and avoid all obstacles along the tour more actively. “Guiding” mode is only entered if the patients have lost orientation to guide them back to their room. As mentioned above, this shall address the patients’ insecurity and anxiety (“Will I find my way back?”) which is a possible reasons for not performing self-training.

During the user trials, our partners from social sciences (SIBIS Institute Berlin) observe and evaluate the training procedure. However, for the later practical use of the robot coach in daily clinical routine in the near future, it is not planed to have clinical staff members being involved. The robot will accompany the patients and document the training session and the achieved results (length of the walking tour, duration of training, average walking speed, number of resting breaks, number of wrong decisions, etc.) autonomously. However, the robot is not supposed to suggest modifications of the training procedure, as this can be done only by the doctor in charge. But the recorded training data provides an objective, quantitative basis for adapting the patient’s training procedure to the measurable progress in the rehabilitation.

2.2 Specifics and challenges of the clinical setting

Our test site for doing this walking training, the “m&i Fachklinik” rehabilitation center in Bad Liebenstein, is a complex U-shaped environment (Fig. 3) which accommodates more than 400 patients. The building has eight-floors, so that the robot coach must be able to navigate across different floor levels. However, the autonomous elevator usage could not yet be integrated into the training application so far due to warranty regulations of the elevator producer.

Moreover, the operational environment is highly dynamic and often very crowded. Staff working in the patients' rooms and patients are moving in the corridors and in the public areas, many of them using walking aids (wheeled walkers, wheel-chairs, crutches) which makes person detection and re-identification very challenging. Often beds, supply and cleaning carts, or wheel-chairs are occupying the hallways, resulting in very restricted space conditions at some times (see Fig. 1). All this requires situation-aware and polite navigation and interaction abilities to guarantee a successful and joyful walking training that will be accepted by both patients and staff.

2.3 Technical requirements for the walking coach

Based on the described training procedure and the specifics and challenges of the clinical setting, the main technical requirements we obtained from the requirement specification with medical and physiotherapeutic experts in clinical post-stroke rehabilitation as well as our own experiences from former assistive robotics projects are summarized subsequently.

An intuitive and robust patient-robot interaction plays a central role in our scenario, as the patients are to be motivated repeatedly using the robot coach during their self-training. Therefore, for *(I)nteraction* between robot and patient, the following requirements have been defined as mandatory:

I1: to reliably detect and keep track of moving, standing, or sitting persons in the local surroundings of the robot even under hard conditions, for example, if patients use walking aids or sit in wheel-chairs,
I2: to autonomously orient towards the current user or drive in a position facing the user as prerequisite for GUI-based interaction,
I3: to robustly re-identify the current user to avoid too many mis-matches and training breaks if the user was temporarily out of view or occluded by other persons or obstacles,
I4: to follow the user in adequate distance during the training,
I5: to guide the user during the training through the center,
I6: to express interaction interest by controling the viewing direction of the robot's eyes in the robot head, and
I7: to realize an intuitively understandable multi-modal (GUI, touch, speech synthesis) dialog for getting and staying in contact with the user (see Sect. 4).

For autonomous, human- and situation-aware *(N)avigation* of the walking coach, the following requirements were specified as crucial:

N1: to allow for a simple and quick mapping of the operational area in the rehab center during the installation of the robot by manually driving the robot around,
N2: to guarantee a robust and precise autonomous self-localization of the robot at all levels of the rehab center,
N3: to drive to any given destination in the center,
N4: to reliably avoid collisions with all possible static and dynamic obstacles in the operational area,
N5: to politely pass standing or walking people guaranteeing a socially acceptable navigation,
N6: to predict and evaluate forthcoming critical deadlock situations and react proactively, e.g. by waiting in an undisturbing position in front of the bottleneck and leaving oncoming persons pass by, and
N7: to autonomously drive and dock to the charging stations in the center.

Taking the training scenario and the requirements and constraints of the clinical setting into account, the following section briefly discusses related work in the field of rehabilitation robotics with a focus on socially assistive robots for walking and orientation training.

3 Related work

A comprehensive overview of current robotic technology in rehabilitation care is given in (Andrade et al. 2014), (Wade et al. 2011) and (EU-Robotics 2015). According to that, up to now the common approach in the field of rehabilitation robotics is the application of orthoses – robotic solutions that physically interact with persons with motor deficits. This includes lower extremity devices such as the LOKOMAT and ALEX (Active Leg EXoskeleton) and upper extremity devices that measure and apply forces and torques to the patient's arm to assess or encourage specific motor task practice. A systematic review of studies and publications dealing with robot-mediated upper limb rehabilitation in stroke is presented in (Basteris et al. 2014). Intelligent walkers, so-called smart walkers or iWalkers (Rodriguez-Losada et al. 2005), (Hirata et al. 2007), equipped with navigation and guiding capabilities also have some bearing to ROREAS, as they try to assist elderly or disabled people in walking alone using the active physical motion support of the walker.

These works in rehabilitation robotics are, however, less relevant for our approach, as ROREAS belongs to the field of *Socially Assistive Robotics* (SAR). SAR is defined in (Wade et al. 2011) as "provision of assistance through social (not physical) interactions with robots. A SAR system uses noncontact feedback, coaching, and encouragement to guide a user during the performance of a task. SAR systems can demonstrate task goals, monitor the user, and provide augmented performance feedback". Although SAR has shown promise in a number of domains, including skill training, daily life assistance, and physical therapy (Feil-Seifer and Mataric 2011), there is no SAR project known that aims in the same direction as ROREAS - the development of a mobile robot coach which can accompany inpatients fully autonomously during their walking and orientation training within a clinical setting.

Therefore, the so called *tour guide robots* are at least of a certain relevance for ROREAS. Among them, there are such well known robots as *Rhino*, *Minerva*, and *Sage*, the exposition guide *RoboX*, or the robots *Mona/Oskar* at the Opel sales center at Berlin (see (Jensen et al. 2005) for an overview). Usually, all these robots guide visitors to a set of exhibits while offering related information, and thus show some similarity to the walking coach function in ROREAS. The same also applies to the still relatively small group of robotic shopping assistants, such as *RoboCart* (Kulyukin et al. 2005), the *ShopBot* robot (Gross et al. 2008), and its successor *TOOMAS* (Gross et al. 2009). These shopping guide robots contact potentially interested customers within the stores and offer their main service, namely to guide the customers on the shortest possible route to the goods shelves with the wanted products. Of similar relevance are the *Zuse-Guide* project (Stricker et al. 2012), where a robot-based mobile visitor information system guides visitors to labs and offices in a crowded multi-level university building, or the *SPENCER*-project (Triebel et al. 2015) where a socially compliant mobile robot, that can assist, inform, and guide passengers in large and busy airports, is being developed. In this context, socially acceptable collision avoidance for mobile robots that navigate among pedestrians is a hot topic, e.g. in (Shiomi et al. 2014).

What all these robots have in common is the need for solving hard technological challenges, as for example politely moving in a crowded environment, guiding interested people and distinguishing them from bystanders, and robustly localizing the robot within this environment. In contrast to the ROREAS-project, all these robot guides only have an instrumental function - namely guiding an interested user on a pre-defined tour or on the shortest possible route to a target position. In these application scenarios, the robot guide typically takes the initiative, and the user has to follow the robot more or less strictly with goodwill. In ROREAS, however, the patient takes the initiative and decides how and how long the walking training has to proceed, whereas the robot has to accompany and actively observe the patient and the training process to assist in appropriate manner if necessary. For such a robot coach that is supposed to train with the same patient again and again for a few weeks, besides the pure instrumental training function a patient-robot relationship needs to be established. In this relationship such social-emotional factors, like the co-experience (how individuals develop their personal experience based on social interaction with a robot), safety (the feeling of security when interacting with a robot), and joy of use (the perceived enjoyment when interacting with a robot) (Weiss et al. 2011) will play an important role for the acceptance and success of a robot-assisted training.

At the functional level of real-world navigation and human-robot interaction, some similarities also exist to the so-called care robots supporting the independent living at home, developed for example in the German projects *WiMi-Care* (Jacobs and Graf 2012) or *SERROGA* (Gross et al. 2015), or in the EU-FP7 projects *CompanionAble* (Gross et al. 2012), (Schroeter et al. 2013) or *Hobbit* (Fischinger et al. 2014). However, none of these care robots was or is involved in such challenging training tasks with disabled people as required in ROREAS. What is still lacking in all these applications of assistive robotics is a strongly human-aware, polite and attentive social navigation and interaction behavior as it is necessary for a rehab assistant that can motivate patients to start, continue, and regularly repeat their self-training with joy.

4 Robot platform ROREAS

For realization of the functionalities and requirements defined in Sect. 2, the technology of the robot platform requires high performance computational units for the execution of all interaction, navigation and service algorithms often running in parallel, intuitive interfaces adequate for disabled people, and multiple sensor systems to perceive the robot's environment and the user. Moreover, the system design had to consider that the robot typically needs to move in a narrow and populated clinical environment. As a consequence, numerous requirements to the design, the technical realization, and the sensor equipment of the robot platform were derived which have directly influenced the design process and the functionality of the robot assistant. In

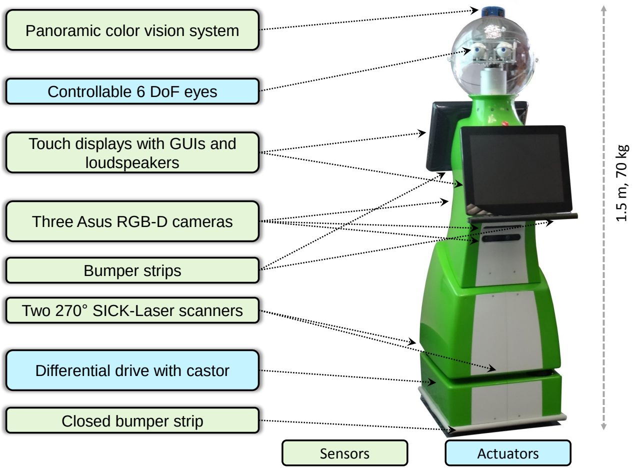

Fig. 4 Robot platform ROREAS developed as part of the project on the basis of a SCITOS [saitoz] G3 robot of MetraLabs Robotics Ilmenau with its main equipment for environment perception, navigation, and HRI.

addition to these functionalities and a pleasant design, later production and operating costs and the longevity of system components had to be considered.

According to the specific requirements for a robotic walking coach, an appropriate training platform has been developed in the project (Fig. 4). Its relatively small size of 45 x 55 cm footprint and a height of 1.5 meters is optimized for a friendly appearance and an ergonomic operation even under limited space conditions.

Drive system: The drive system of the robot is a differential drive with a castor on the rear which gives the robot a good maneuverability and stability in spite of its size and weight of 70 kg and allows a maximum driving speed of up to 1.2 m/s.

Sensor equipment: For user perception, human-aware navigation, and collision avoidance, the robot is equipped with an innovative sensor concept consisting of two SICK laser range finders with a total scanning range of 360^o placed at a height of 20 cm above the floor, three Asus RGB-D cameras (two in driving direction, one backwards), and a panoramic color vision system mounted on top of the head.

Human-robot interaction: For interaction with the patients, ROREAS is equipped with two 14 inches touch displays for use while standing as well as sitting, a sound system, and a robot head with two eyes with six degrees of freedom (DoF). The touch displays are the central communication interfaces to the robot, and the head gives the robot a smart but still technical appearance, which encourages users to interact with it. The head has the following degrees of freedom: lifting and lowering of the head $(+20^o, -7^o)$, rotation of the whole head (350^o), synchronized up and down movement of the eyes, synchronized left and right movement of the eyes, opening and closing of the eye lids (independently for each lid). This way, the head directly supports the user-robot interaction by the controllable viewing direction of the eyes. This eye contact has proven to be very helpful for successful interaction and exercising.

Power capabilities and operating time: With all the sensors and actors running, a hierarchical energy-saving concept in conjunction with energy-saving units enables a long run-time of about 8 hours until the robot needs a break for recharging. Assuming a maximum duration of one hour for a training session, the robot can accompany eight patients over a full work shift without breaks. It can be recharged by the integrated charging system in about 6 hours. As the robot can autonomously dock to its self-charging station, a 24/7 operation is possible if the resting periods of the patients are strictly used for recharging.

Safety mechanisms: In addition to the laser- and vision-based 2D and 3D obstacle detection (see above) and the reactive navigation skill used for obstacle avoidance and navigation safety described in Sect.6.5, the robot is equipped with multiple bumpers, to detect possible physical collisions with obstacles. This additional safety system involves a closed bumper strip about 4 cm above the ground, and bumper strips at both touch displays. All bumpers are directly coupled to the drive unit and guarantee an immediate safety stop, whenever a collision is registered. For safety reasons, continuing the tour requires an explicit confirmation by the user. For an early detection of stairs going down, the robot is additionally equipped with a distance measuring *TeraRanger One* ToF-sensor which is 45^o tilted in forward direction and placed beneath the touch screen.

In comparison to the preliminary experimental platform presented in (Gross et al. 2014), this new robot platform is explicitly tailored to the user group of stroke patients with a focus on easy usability while standing or sitting, joy of use, and positive user experience, but also on later production and operational costs.

5 Functional system architecture

The functional system architecture of the robot coach is a further development of the architectures of our shopping guide robots (Gross et al. 2008) (Gross et al. 2009). In comparison with these, however, the ROREAS architecture is more complex, includes more human- and situation-aware navigation and interaction skills and behaviors (Fig. 5), and allows more flexibility in realization of future new applications (training functions). In the ROREAS architecture, all robotics related methods and skills have been consistently abstracted

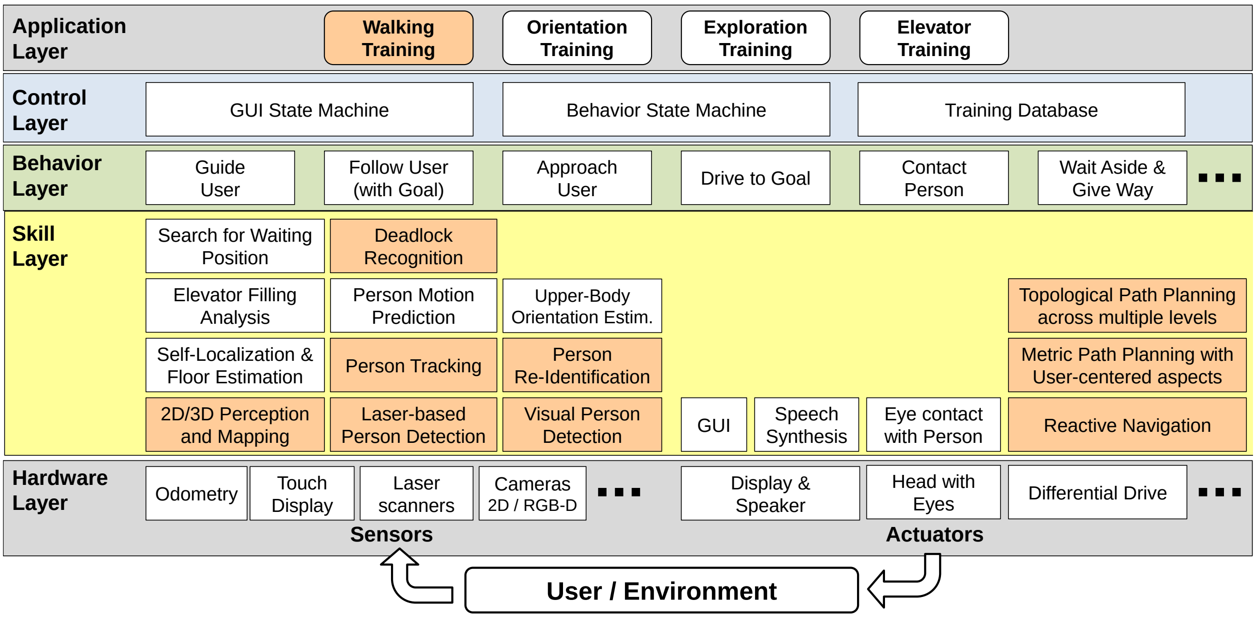

Fig. 5 Multi-layered functional system architecture of the ROREAS training assistant consisting of *Hardware Layer*, *Skill Layer* with navigation and HRI-specific methods and skills, *Behavior Layer*, *Control Layer* orchestrating the behaviors, and *Application Layer* implementing the specified applications for post-stroke self-training. Only the reddishly highlighted modules will be covered in this paper as they are of particular relevance for a human- and situation-aware navigation of the robot coach.

from the applications itself resulting in a flexible, layered system architecture. The bottommost *Hardware Layer* encloses the hardware (sensors and actuators) and the operating system. The low-level sensor information is processed in the next higher level, the *Skill Layer*, which covers the whole spectrum of required navigation and HRI skills that are executed in the *Hardware Layer* in parallel.

Above the skills there are diverse modules representing the *Behavior Layer* which make use of the HRI and navigation skills in the layer below. Here, for example, a “Guide user” behavior is realized, other behaviors are “Approach user”, “Follow user”, or “Wait aside” which are necessary for direct interaction as well as a polite human-aware navigation. These behaviors are exclusive units each representing an individual control loop for accomplishing the different navigation and interaction functions of the robot. To do so, the currently used behavior activates, deactivates, and parametrizes the required skills. Furthermore, the *Behavior layer* operates as an interface for the *Control Layer* where two finite state machines, the *GUI-* and the *Behavior State Machine*, and the *Training Database* are implemented. This layer contains the behavior control, which makes use of the basic features provided by the HRI and navigation skills and is orchestrating the behaviors. Based on the *Training Database* and closely coupled to the Graphical User Interface (GUI) and the Speech Synthesis, the *GUI State Machine* is responsible for the patient-specific training process taking into account personalized therapy plans and the already achieved progress in self-training. The *Behavior State Machine* comprises a set of states where each state is associated with one of the behaviors in the *Behavior Layer*. Transitions between the states are triggered by navigation events, person tracking events, GUI interaction, or via the administration remote interface. The highest layer, the *Application Layer*, implements the specified applications and services, the “Walking coach” as the currently most important training functionality, and leaves room for further applications, such as the “Orientation training”, the “Exploration training” or “Elevator training”.

The robot’s basic functionalities for user tracking, navigation, and interaction are implemented using MIRA, a middle-ware developed for robotic applications, providing a framework suited to the requirements of distributed real-time software. For an introduction to MIRA and comparison to the popular robotics software framework ROS (Quigley et al. 2009), see (Einhorn et al. 2012).

6 Human- and situation-aware navigation and interaction skills

Since a complete description of all services, behaviors, and skills of the robot coach required for the walking and orientation training would go beyond the scope of this paper, subsequently only an overview of those skills and behaviors is given that are of direct relevance for a human- and situation-aware navigation and training to meet the technical requirements for navigation and interaction defined in Sect. 2.3.

6.1 Person detection and tracking skills

In order to guarantee a successful walking training, at any time the robot needs to know the exact position of its current training partner and other people (staff, patients, visitors, other bystanders) standing, sitting, or walking around in its vicinity (requirement I1). For this purpose, we utilize a probabilistic multi-hypotheses people detection and tracking system (Fig. 5, Skill layer) developed in our lab (Volkhardt et al. 2013). This system is able to track walking people and people in standing or sitting poses. It is based on a 7D Kalman filter that tracks the position, velocity, and upper body orientation of the respective persons. The tracker processes the detections of different, asynchronously working observation modules: a 2D laser-based leg detector, a face detector, and an upper-body shape detector (Weinrich et al. 2012). The laser-based person detection is well suited for detecting pairs of legs as indicator for the presence of people in the vicinity of the robot. However, in a rehab center we have to deal with stroke patients who often need walking aids. These tools occlude or touch the legs of the patients. Therefore, we have advanced the aforementioned leg detector by introducing generic distance-invariant laser-scan features that are then utilized to train classifiers for detecting people without walking aids, people with walkers, people in wheelchairs, and people with crutches (requirement I1). Using this new approach for laser-based people detection, in comparison to the first approach based on (Arras et al. 2007) we could significantly improve the detection of persons with and without walking aids, and now even can classify the different walking aids with 86% correct classification rate (Weinrich et al. 2014).

6.2 Person re-identification skill

As the self-training is to be performed in the corridors of the rehab center, often many other people will be present in the surroundings of the robot (see Fig. 6). To hold contact with the current user, the robot must continuously track the user, and when the person was temporarily lost from view, it must be able to re-identify the user by its visual appearance (requirement I3). However, re-identification of a person by a mobile robot

Fig. 6 The *Person following* behavior needs an efficient and reliable *Re-identification skill* to succeed during rush-hour times in a clinical environment.

is very challenging due to the real-time requirements, motion blur in the video data, a very dynamical environment with many different lighting conditions, and many objects, that trick state-of-the-art visual person detectors to false positive detections. The narrow floors of the building often lead to partial and temporally full occlusions of the user by other people. The user may also be only partially visible, if standing near the robot, due to the camera mounted on top of the robot's head. Additionally, the size of the image regions containing the user varies a lot with the distance to the robot. Therefore, the person re-identification skill (see Fig. 5, Skill layer) has to be robust to image motion blur, varying resolution and illumination, occlusions, people with walking aids, and false positive detections. Moreover, the model of the current user needs to be learned very effectively and quickly (see Sect. 2.2) while the user is standing in front of the robot during logging in via the robot's touch screen (maximum one second for model training). When re-identification is necessary during the training, the current user should be correctly recognized in at least 95% of all cases to avoid too many mis-matches and training breaks (requirement I3). This target number is based on the assumption, that system failures (the number of misidentifications) can be modelled by a Poisson process, where the time-to-failure is exponentially distributed. With 95% correct re-identification rate, the probability to complete a 250 meters long training tour without any person mismatches is 84%, so a typical training tour with a length of only 100 meters could be completed without breaks. A more detailed explanation for this estimate is given in Sect. 7.1.3. When a decision between two nearby person hypotheses is hard to make, the robot should stop and ask the user to make itself felt. When it even loses contact to the user and cannot re-detect its user again, the robot first drives to the nearest resting place to check

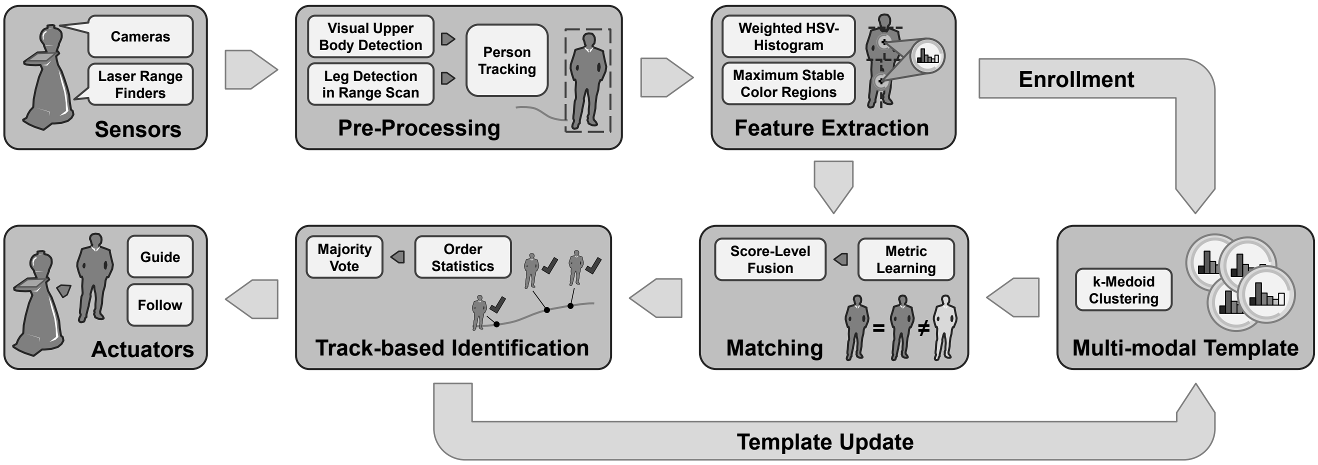

Fig. 7 Overview of our appearance-based person re-identification module running in real-time on the robot (from (Eisenbach et al. 2015a)).

if it can find its user there. If this is without success, it drives back to the patient's room and waits there for a few minutes.

To meet these requirements, we have developed a re-identification workflow that is optimized regarding processing speed, but does not decrease recognition accuracy. Fig. 7 gives a coarse overview: First, all persons in the image have to be detected and tracked (see above). Then, their appearance is described by multiple complementary features. The current user is represented by a multi-modal template composed of features extracted during the log-in at the beginning of the training. To reduce the size of the user template, similar appearances are fused by clustering. To accurately compare all persons in the scene with the template of the current user, we apply a distance metric that has been trained on a scenario-specific dataset. To compose the matching results for the different features, information fusion at score-level (Ross and Nandakumar 2009), (Eisenbach et al. 2015b) is performed. Afterwards, the person hypothesis is chosen by a track-based decision considering preceding observations. Additionally, if the user can be identified securely, the template is updated. Implementation details and experimental results of this re-identification module on diverse standard benchmark datasets are presented in (Eisenbach et al. 2015a).

6.3 2D/3D perception and mapping skill

In complex environments, such as the described rehabilitation center, many obstacles are very hard to recognize by the robot's 2D laser-range-finders, mostly since the main extent of the obstacles is located above or below the plane that is covered by the laser scanner. To reliably avoid collisions with all possible static and dynamic obstacles in the corridors and public spaces (requirement N4), we additionally use three Asus RGB-D cameras to obtain 3D information about the structure of the local surroundings. However, these sensors produce a huge amount of data, and hence an appropriate representation is needed for processing this data efficiently. We use a map representation that is based on the Normal Distribution Transform (NDT) (Magnusson et al. 2009). Such maps achieve a significantly higher accuracy than voxel maps when the same cell resolution is used (Stoyanov et al. 2011). The high accuracy is necessary for a precise navigation especially in narrow corridors. As described in (Einhorn and Gross 2015) in more detail, our NDT mapping approach is able to model free-space measurements explicitly. Moreover, it detects and handles dynamic objects such as moving persons directly within the generated maps. This enables its usage in highly dynamic environments. In addition to the 2D laser-based oocupancy grid maps, the generated local 3D-NDT maps are used for local navigation and obstacle avoidance (see Sect. 6.5).

6.4 Deadlock recognition skill

Due to the structure of the building with long relatively narrow corridors which are often still further narrowed by medical trolleys, stretchers, or wheel-chairs, the robot is often confronted with narrow passages which permit movements only in one direction at a time. Moving in such a confined space imposes *(i)* deadlocks in narrow passages caused by a forthcoming person and *(ii)* the problem of queuing, when the robot and a person are attending the narrow passage in the same direction. Fig. 8 schematically illustrates both situations. Since a polite and proactive navigation is an important

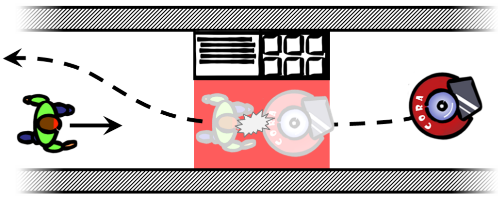

(a) Deadlock situation with a forthcoming person

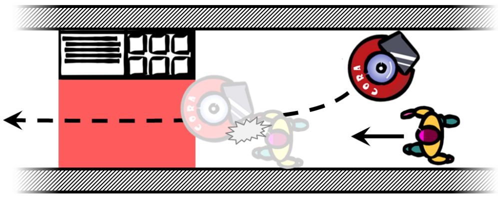

(b) Queuing situation with a person moving in the same direction as the robot.

Fig. 8 Schematic depiction of conflicting situations caused by a narrow passage typical for the clinical setting.

requirement for a mobile robot assistant (requirements N5 and N6), these situations must be recognized in advance to trigger a proactive reaction of the robot, that is *(i)* driving to an undisturbing waiting position to give way for the forthcoming person or *(ii)* forming a queue by following the person through the narrow passage.

In our approach, these situations are characterized by a set of features describing a narrow passage in conjunction with a predicted space conflict with a moving person and the spatial relationship between the person, the robot, and the narrow passage. A possible space conflict is detected by predicting the trajectories for all tracked persons and for the robot through the narrow passage. A collision can be described as that point in time and space, where the predicted trajectories intersect. Narrow passages are detected by calculating the lateral distances to obstacles along the planned path of the robot. The obstacles are extracted from the local 2D/3D navigation maps (see Sect. 6.5). Along with the predicted collisions of all tracked persons, features describing the spatial situation around the narrow passage are extracted. To provide the required polite behaviors at narrow passages, currently we distinguish three different classes: *Waiting*, *Queuing*, and *Proceed* for non-conflicting situations. As situation classifier, we use a hand-designed decision tree. If class *Waiting* is recognized, the robot rapidly needs to find an appropriate undisturbing waiting position in front of the deadlock. For this purpose, we use a combination of cost functions that assess the suitability of possible positions with regard to a set of specific criteria: *(i)* to allow the robot to wait near walls, *(ii)* not to obstruct moving people, *(iii)* to place the robot as close as possible to its current position, and (iv) to allow the robot to continue observing the narrow passage. By using a Particle Swarm Optimization (PSO) approach (Kennedy 1995), the pose with minimum costs is used as temporary waiting position. A detailed description of the *Deadlock recognition skill* is given in **?**.

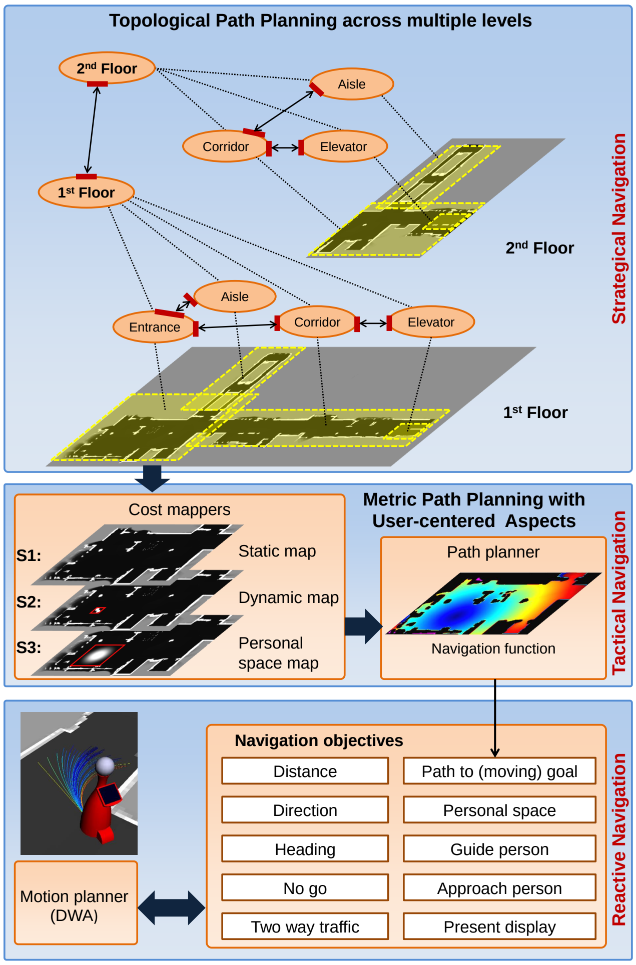

Fig. 9 Detailed view of the navigation skills shown in Fig. 5 right. The *Topological Path Planner* (Strategical navigation) and the *Metric Path Planner* (Tactical navigation) inject global knowledge into the *Reactive Navigation* via the navigation objective *"Path to (moving) goal"*.

6.5 Navigation and obstacle avoidance skills

In addition to a reliable and intuitive human-robot interaction, robust and user-centered navigation is a fundamental requirement for an autonomous robot coach. The core components of our navigation architecture can be classified in reactive, tactical (based on metric path planning), and strategical (based on topological path planning) navigation located within the *Skill Layer* (see Fig. 5, right). These components are shown more detailed in Fig. 9.

On the reactive level, the DWA-based motion planner (Fox et al. 1997) and so-called navigation objectives determine motion commands for the robot's motor controllers according to the current task, which is set by the current behavior localized in the *Behavior Layer* (see Fig. 5). Each navigation objective is a separate software plugin specialized for a certain navigation sub-task, such as following a path to a goal position, following the current user, approaching a standing or sitting person, respecting personal space, or avoiding obstacles by means of the robot's *2D/3D perception and mapping* skill (see Sect. 6.3). This allows us to add new objectives easily, when new tasks and behaviors become necessary without changing existing parts of the navigator. A detailed overview of all implemented objectives is given in (Gross et al. 2014).

The output of the objectives is then used by the *DWA-based Motion Planner* (Fig. 9, bottom left) to generate motion commands that are then sent to the robot's motor controllers. For evaluation of all possible motion commands within the DWA, the objectives require additional information from other modules of the *Skill Layer*, such as person hypotheses from the person tracker (see Fig. 5, left). Therefore, the objectives can access this information directly from these modules. To rate all possible motion commands and search for the currently best suited one, most objectives use short-term predictions of the robot's clothoid-shaped local motion trajectories (Fig. 9, bottom left) in combination with the cost maps (Fig. 9, middle left) which are based on 2D occupancy grid maps. In these occupancy maps all static obstacles (walls, doors, etc.) and dynamic obstacles (persons, wheel chairs, supply and cleaning carts, etc.) captured by the robot's sensor system (see Sect. 4) are represented. To include the knowledge about static and dynamic obstacles represented in the local 3D-NDT map (see Sect. 6.3) into this motion planning, the local 3D-NDT map is simply projected down to the 2D plane and merged with the laser-based 2D occupancy map.

To allow path planning across multiple floors and to decrease the computing effort for path planning on metrical maps, we utilize a hybrid, hierarchical topological map for path planning at the strategic level (see Fig. 9, top), which allows us to model the elevators as transitions between the different floors of the rehab center. On the coarsest level of this graph, each node represents a single floor of the building. Each node is further subdivided into sub-nodes that represent the aisles of each floor, etc.. On the finest level, the leaf nodes contain metric occupancy maps. The path planning starts on the coarsest level using a general Dijkstra algorithm and is iteratively refined up to the metric occupancy maps (Fig. 9, middle), where we finally apply the computationally more expensive E* path planning algorithm (Philippsen and Siegwart 2005). This hierarchical approach combines fast path planning on a topological map and the ability of dynamic re-planning that is supported by the E* algorithm.

7 Field tests of the walking coach

Before it was possible to evaluate the walking coach together with stroke patients in user trials, it had to be assured that all the required skills and behaviors for HRI and human-aware navigation (see Fig. 5) did work as expected in the clinic setting. Therefore, our field tests were tripartite: First, we evaluated all skills and behaviors in a controlled setting in the operational environment of the rehab center under everyday conditions. We began with the evaluation of the navigation skills. Afterwards, we added skills and behaviors that included interaction with a user. For security reasons, we did not interact with actual patients at this stage, but replaced them by briefed technical staff of our robotics lab as described in Sect. 7.1. When the functional field tests had been completed successfully, in the second stage we tested the walking coach with clinical staff imitating the walking behavior of stroke patients, to ensure a correct behavior of the robot coach in all defined situations (Sect. 7.2). Based on this, when the walking coach had proven to be reliable and suitable, we started with first user trails with actual stroke patients, which are described and analyzed in Sect. 7.3.

7.1 Functional field tests with project staff members

To ensure, that all skills and behaviors required by the walking coach do work accurately and securely, we at first performed functional on-site field tests with staff members of our robotics lab. This required the definition of benchmarking scenarios and systematic studies (EU-Robotics 2015). To this end, the requirements defined in Sect. 2.3 were assessed in these functional tests. With respect to the *navigation functionality*, the requirements N1, N4, N5 and N7 could be fulfilled successfully from qualitative point of view. Only requirements N3 (drive to any given destination in the center), N2 (self-localizing at all levels), and N6 (handling deadlock situations) are still ongoing issues that need to get solved the elevator control problem and the early detection of deadlock situations. However, this was not an impediment to the user trials, since the walking training (in comparison to orientation training) will only take

Tested Skills and Navigation Objectives	**Driving** (N3)	**Guiding** (I5)	**Following** (I4)
2D Obstacle Avoidance	D1	-	-
3D Obstacle Avoidance	D2	-	F1
Personal Space	D3	-	-
Deadlock Recognition	D4	G1	-
Driving on the Right	D5	-	-
Person Re-Identification	-	G2	F2

Table 1 Overview of the conducted tests for the assessment of the basic behaviors and use cases "Driving to specific goals" (N3), "Guiding a person"(I5), and "Following a person" (I4) together with the skills and objectives activated in each test case.

place on a single floor, and handling of deadlock situations is necessary in guiding mode only. In the *Walking coach* user trials with patients, the robot shall follow them most of the time and guide them in rare cases of emergencies only (e.g. when patients loose their bearings).

A detailed quantitative analysis of the robot's navigation capabilities regarding obstacle avoidance, polite distances to humans, the utility of considering the user's personal space, or the recognition of forthcoming deadlock situations will be presented subsequently. Regarding the required *HRI functionality*, the situation is very similar - most of the use cases have already been tested on-site with volunteers and could be demonstrated successfully from qualitative point of view. Additionally, first quantitative results of benchmarking the *HRI functionality* are presented subsequently.

7.1.1 Test design and realization

Extensive functional testing was performed in February 2015 over the course of 4 days and a driven distance of 15,000 meters within several floors of the "m&i Fachklinik" rehab center in Bad Liebenstein at different times throughout the day. This was done to assess the robot's basic behaviors under varying conditions, such as challenging building-structures, changes in illumination, and a variable amount of people within the corridors. Table 1 provides an overview of the basic behaviors that were evaluated during functional testing, namely "Autonomous driving to specific goals" (N3), "Guiding a person"(I5), and "Following a person" (I4) together with the skills and objectives (see Fig. 5 and 9) used in each case (e.g. 2D vs. 3D obstacle avoidance).

The general setup for the functional tests was as follows: The robot had to drive autonomously back and forth between navigation points that were located at both ends of three hospital floors (see Fig. 10), with an activated skill set depending on the current test case (see Table 1). For quantitative assessment of these skills, measures, as e.g. the number of collisions or person mismatches, or the needed travel time, were determined. Additionally, an external observer was present, who accompanied the robot and documented its behavior, but always from far distance to prevent any distraction. Regarding the *navigation performance*, the observer counted the number of (i) close (less than 10-15 cm) passings of obstacles, (ii) close passings of persons, and (iii) manually triggered emergency stops. Regarding *deadlock recognition at narrow passages* (analyzed separately for people moving in the same and in the opposite direction compared to the robot's current course), it was counted if a deadlock situation was correctly recognized by the robot (true positives), ignored (false negatives), or erroneously detected (false positives). Regarding *person recognition during guiding and following*, it was determined, whether and how often the robot confused the current user with a different close-by standing person.

For the sake of an interruption-free testing process, the external observer used a control interface on a tablet-computer (called *control tablet*), which was specifically developed for this purpose. Using this handy tool during the field tests of the following behavior, the observer was able to make real-time adjustments to skills (such as person detection and recognition) and compensate erroneous decisions of those skills which are still under development. This way, the functional tests in the real clinic environment could be started much earlier than this would have been possible from the readyness level of the respective skills. Moreover, the developers got objective and situation-specific feedback about the functioning of their algorithms. At the same time, each manual intervention (type and quantity) was documented and stored within the device and could be condensed into a graphical event-log afterwards.

7.1.2 Evaluation of navigation skills

2D vs. 3D Obstacle avoidance (D1 vs. D2): In extensive field studies, we evaluated the performance of our *3D perception and mapping* skill (see Sect. 6.3) in comparison with a standard 2D approach to meet requirement N4 at best. In these tests conducted on three different floors of the rehab center, the robot was driving autonomously between two navigation points that were located at both ends of the floors (see Fig. 10). The traveling distance between the two goals was 172 m. The robot's total mileage in these tests was 12,000 m. Hence, the robot commuted between both goals approx. 70 times. The average driving speed was 0.6 m/s, the maximum 1.0 m/s, which is sufficient as the patients' typical walking speed is significantly lower. To evaluate the obstacle avoidance capabilities, we counted the collisions and close encounters (less than 15 cm) with obstacles that nearly resulted in collisions. The latter were identified and tagged manually via the control tablet by the observer who supervised the robot's behavior. The obstacles that had to be avoided between the two goals were three medical trolleys, three seating-accommodations, two or more wheel chairs, and dependent on daytime, up to three supply carts of cleaning staff and up to 30 people (patients, clinical staff, visitors) as dynamic obstacles.

During the first 2,000 m of these tests (D1), the 3D perception was *disabled*, and the robot sensed its local surroundings solely using its 2D laser range finders. In these tests, we counted 23 collisions and 18 near-collisions, which on average correspond to 12 collisions and 9.4 near-collisions per kilometer. Due to the high risk of colliding with obstacles or patients, we aborted these tests after 2,000 m. During the remaining 10,000 m of our test runs (D2), we activated the 3D perception. As a result, the number of collisions could be dramatically reduced. There were only one collision and 4 near-collisions, which on average correspond to 0.1 collisions and 0.4 near-collisions per kilometer. Having measured these mean failure rates, we could conduct a reliability analysis in order to estimate the probability for finishing a training successfully without collisions. The mean traveling distance for the robot to perform a full training with a patient can be expected as $s = 250\ m$, including the way to the room of the patient and the training itself. Just like other system failures, the number of collisions can be modelled by a Poisson process, where the time-to-failure (here the driving distance to a collision) is exponentially distributed. Hence, the probability for driving a distance s without collision is given by $p(x > s) = e^{-\lambda s}$ where λ denotes the failure or collision rate (Balakrishnan and Basu 1996). Using the measured collision rates, the probability to complete a 250 m training tour without any collision is 97% if 3D perception is activated. If no 3D perception was used, the chance for a successful training decreased to only 5%. It should be noted, that for our 3D perception and obstacle avoidance approach the real-time requirements were always met, while the module only used 5-10% of the on-board computational power of the i7-CPU on average.

Influence of considering personal space (D3): To evaluate the capabilities of our approach for politely respecting the personal space of bystanders while navigating to a goal (requirements I4 and N5), a similar setup to the 3D obstacle avoidance evaluation was used (see Table 2). In these tests, we used the same goals on

	3D perception only	**3D percept. + Person. space**
Mileage	1,500 m	8,000 m
Persons	23 (15 per km)	133 (17 per km)
Near encounters (less than 15 cm)	11 (7 per km)	4 (0.5 per km)

Table 2 Results from the evaluation of the *Personal Space* objective (see Fig. 9).

the three different floors. The complete evaluation covered a total mileage of 9,500 m. The evaluation was split into two runs again. During the first 1,500 m, the robot ignored the personal space, while during the remaining 8,000 m the personal space was respected. Due to the convincing result of the 3D perception and obstacle avoidance, we enabled the 3D perception for both test runs. For evaluating the capabilities, again an observer counted critical near encounters with persons, i.e. when the robot drove too close to a person (less than 15 cm), as this causes discomfort by confusing the person about the robot's interaction intention or even violating the person's privacy. Since this measure strongly depends on the traffic volume, we conducted the tests at the same time of subsequent days. To compare the traffic volume, the overall number of persons who passed the robot in a radius of 2 m was counted. Both test runs had a comparable traffic volume. Regarding the near person encounters, the first run ignoring the personal space had 11 encounters with an average of 7 near encounters per kilometer. During the second run with the personal space enabled, only 4 encounters were counted, resulting in an average of 0.5 encounters per kilometer. From these results, we conclude the necessity of a navigation behavior considering the personal space for the clinical environment.

Polite behavior at narrow passages: In this test, we evaluated the performance of the *Deadlock recognition* skill relevant for requirement N6. For minimizing the risk of collisions or disturbance of patients, we enabled the *3D perception* skill and *Personal space* navigation objective (see Sect. 6.5) for this test. During the test, we let the robot autonomously drive between different goals or guide a single person to these goals. In total a distance of 4,700 m was traveled, and 157 bystanders (staff members, patients) were crossing the robot's way in a 2 m radius. From these 157 bystanders, 35 persons caused potential deadlock situations (15 queuing and 20 waiting situations) at narrow passages on the corridor the robot had to react on. The test conductor manually counted the decisions taken by the *Deadlock recognition* skill. The results are as follows: 34 of 35 deadlock situations were correctly classified (true positive rate of 97%), only one deadlock situation was not recognized by the robot. From these 34 correctly classified deadlock situations, 30 were assigned to the correct deadlock type (accuracy of 88%). However, there was a not negligible number of false positives - 26 of 122 uncritical situations were classified as deadlocks (false positive rate of 21%) initiating a stopping of the robot and waiting aside. For an everyday suitable walking training with patients, this is still an unacceptable result, as the wrong deadlock decisions again and again would unnecessarily interrupt the training process with negative consequences for the user acceptance. By critically analyzing these results, it revealed that the performance of the *Deadlock recognition* strongly depends on the accuracy of the situation describing features which in turn depends on the person tracker, the narrow passage detection, and the space conflict prediction. Analysis of the false positives revealed that 19 of 26 false positives are caused by false detections of the person tracker. In these cases, the deadlock recognition assumed to have a conflicting situation with a person, even though there was no person present at all. The remaining 7 false positives were caused by dynamic obstacles, e.g. moving persons or objects moved by persons. Since the narrow passage detection uses the navigation map which currently is not yet able to distinguish dynamic obstacles from static ones, this movement could not be considered in the space conflict prediction and leads to false predictions. Thus, for a more robust human-aware situation recognition further improvement of the person tracker and the deadlock recognition are required, which is subject of ongoing work. For the functionality of the *Walking coach*, however, this behavior is not absolutely necessary at the moment, as in this mode the robot is following the patient, and the patient has the initiative and tries to avoid deadlocks.

7.1.3 Evaluation of user re-identification

Test Design: To evaluate the benefit for the robot to utilize the user *Re-identification skill* for resolving ambiguity in user tracking, we also performed live tests in the rehab center. Over a period of six hours, the robot followed and guided three probands through one corridor of the center (Fig. 10). Their appearances cover typical clothing: dark/black, light/gray, and colored clothes. To check, if there is still contact to the current user, the person re-identification module frequently has to compare all persons in the robot's surroundings with the currently valid user model learned during logging in. In each run, one of the probands was guided and followed by the robot (tests G2 & F2 in Table 1) as shown in Fig. 10 for a distance of 400 m. The probands

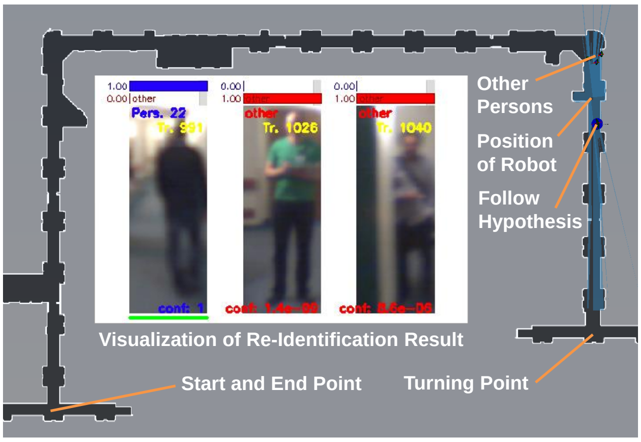

Fig. 10 Map of one floor of the rehabilitation center in Bad Liebenstein. Center: Exemplary visualization of appearance-based re-identification where three probands were standing around the robot.

could freely choose their route and walking speed, but were instructed to behave like stroke patients (i.e. no running). The behavior of the robot was observed and manually corrected via the control tablet whenever the robot did not succeed. In these cases, the correct position of the test person was manually marked within the local map shown on the control tablet, and the robot had to continue. We repeated guiding and following until a pure driving time of one hour was reached for each proband and test. Therefore, the robot guided the probands for an overall distance of 2,000 m and followed the probands for 2,400 m. The behavior of the robot was rated as correct, if it could follow or guide the correct person continuously. Short stops, for a maximum of three seconds, were accepted for situations where the robot was uncertain (e.g. when the person was not detected for several consecutive images). Then it had to continue. If contact to the test person was lost, the robot had to stop, and was not allowed to follow/guide other persons. Otherwise, the stop was forced manually via the control tablet, the correct person hypothesis was reset, and the robot had to continue.

To evaluate, which decision the robot would have made, if it had not used the appearance-based visual re-identification, a reference approach was run simultaneously: whenever the user track broke, it simply did choose the new track with the shortest spatial distance to the last observation. The robot was forced, however, to ignore this distance-based decision and behave like the re-identification module suggested, as only the number of these failure events was of interest. We expected a significant reduction of person switches using visual re-identification in comparison to the simple spatial distance-based approach. For all these tests, the *3D perception*, *Personal space*, and *Deadlock recogni-*

Following the User

Run	Near Persons	Gener. Tracks	Mism. Pers.	Mism. Non-pers.	Unnec. stops	Mism. Pers. Ref.
1	15	748	1	1	0	2
2	13	701	0	0	1	3
3	14	262	1	0	1	1
4	11	772	0	0	0	1
5	8	241	0	1	0	3
6	6	275	1	0	0	0
$\sum$	67	2999	3	2	2	10

Guiding the User

Run	Near Persons	Gener. Tracks	Mism. Pers.	Mism. Non-pers.	Unnec. Stops	Mism. Pers. Ref.
1	8	386	0	1	1	1
2	13	465	0	0	2	1
3	10	847	0	0	0	0
4	5	247	0	0	1	0
5	12	154	0	0	0	1
$\sum$	48	2099	0	1	4	3

.

Table 3 Re-identification performance in live tests in the rehab center. **Near Persons**: number of nearby persons while following the user for a distance of 400 m along the floor of the rehab center. **Gener. Tracks**: number of assigned tracking IDs for new person hypotheses (including all tracks for near and far persons and valid false positive detections). **Mism. Pers.**: number of mismatches of the current user with other persons. **Mism. Non-pers.**: number of mismatches of the current user with non-person objects, not caused by re-identification. **Unnec. Stops**: number of unnecessary stops triggered by the re-identification module, even though the user was visually still observable. **Mism. Pers. Ref.**: number of mismatches of the current user with other persons using the reference method (see text) without visual re-identification.

tion skills were activated in combination with *Person re-identification*.

Table 3 shows the results for following and guiding the user. As can be seen, the implemented person re-identification performs well and helps the robot to decrease the number of mismatches with other persons. A drawback of visual user detection is the occasional confusion of the current user with non-person objects (see Table 3 - *Mism. non-pers.*). This happens if misaligned images showing background structures or false positive detections appeared during the user model learning phase after logging in. Then, the learned user template incorrectly consists of correct user observations but also of erroneous images of other person-like objects in the scene, that may produce false positive matches with these objects later on. In these cases, manual intervention was necessary three times, when the robot followed false positive detections and could not resolve the situation by itself. Since the number of mismatches needs to be further reduced, the accuracy of the used *Person detection skill* has to be improved further to guarantee an error-free user template training after log in. The robot did very well in stopping when the user was temporarily not visible. At rush-hour times, where the reference approach (see Table 3 - *Mism. pers. Ref.*) fails clearly, the visual re-identification performed very well. For example, in the situation shown in Fig. 6, the robot had to follow the proband on a zigzag course through a group of seven people. The user was traced almost through the group, but then he was lost during an evasive maneuver. The robot immediately stopped as desired.

Summarized, during the two hours of following and guiding probands on a track of 4,400 m, the robot came in close contact with 115 other people. Overall, the user was mismatched only three times. Even at rush-hour times, the robot was able to reliably follow and guide probands through the corridor of the rehab center. We can expect 0.6 mismatches with other people per kilometer, or 1.2 mismatches per kilometer in total (additional mismatches with false positive detections). As described in Sect. 7.1.2, the occurrence of mismatches can be modeled as Poisson process. Using the measured mismatch rates, the probability to complete a $250\,m$ training tour without any person mismatches is 84%. Considering also false detections, the probability to complete the training without any re-identification errors is 74%. The current performance is acceptable for first user trials with actual patients, described in the next section, when an observer can correct the rare wrong decisions via the control tablet. For autonomous long-term operation with patients, however, the re-identification performance still needs to be improved further.

7.2 Functional tests with clinical staff

After successfully completing the functional tests with staff members of our robotics lab, in May 2015 we evaluated the walking coach again - but this time with the help of clinical staff. In these trials, trained clinical staff imitated the walking behavior of typical stroke patients, and the robot had to accompany them during their self-training. This included extended human-robot interaction, like interaction via touch screen, acoustic feedback to the user via speech output, and continuous eye contact with the currently tracked user to intuitively demonstrate the current point of interest. As in the first functional tests, the trial observer could correct wrong decisions of the robot's skills via the control tablet. With this support, all skills and behaviors proved to be accurate and secure, so we could begin with the first user trials with actual stroke patients in June 2015.

7.3 User trials with stroke patients

In these user trials, only volunteers from the group of stroke patients who already got the permission for doing self-training by their doctor in charge have been involved. On the day of the testing, these volunteers got a demonstration and training of the robot's functions and abilities. Throughout the whole tests, the users' activities were temporarily logged in specific log-files used in ROREAS for automatic recording of the service usage, unexpected events, and necessary corrections of the external observer via control tablet. During the user trials, an instructor and technical staff was present at all time to guaranty safety for the patients. The technical staff observed the intentions of the robot on the control tablet and intervened early whenever the robot replied that it is unsure about its decisions (e.g. in re-identification) or did not behave as expected. Due to data protection regulations, we only were allowed to store a limited amount of data. Therefore, we can only report first quantitative results.

First user trial in June 2015: The first user study was performed during low traffic times with defined training routes at the ward. The robot performed eleven walking training tours with five different patients (1-3 per patient depending on their fitness) who used different walking aids. The training only took 62 minutes, and a distance of 873 meters was covered. During these tests, the robot collided three times with obstacles. These collisions resulted from rotational movements near walls with handrails that could not be observed by the robot's sensors. Moreover, the robot came in close contact with 78 people. Only twice, the robot violated the personal space of a person. In both cases, this behavior was hard to avoid due to high traffic on the cooridors. Manual intervention by means of the control tablet was performed 19 times, to interact with the re-identification module. Only two of these cases were false decisions, the others were feedback due to uncertainty of the recognition module.

Second user trial in September 2015: After these first user trials with patients, we analyzed failures, improved the skills and behaviors of the robot coach, and successfully re-tested the robot coach with staff members in August 2015. After this, in September 2015 we conducted a second user trial with stroke patients in a more advanced setting, which included longer training routes that could be chosen by the patients and a training even during rush hour times. In this trial, the robot performed 21 walking trainings with seven different persons (2-4 per patient depending on their fitness). The robot followed the patients for 2,109 meters in 142 minutes (including all interactions and short resting breaks). The number of collisions with obstacles was reduced to one. Probably, this was caused by an outage of the 3D sensor that did not recognize a small wheel of a medical trolley. Also the number of violations of personal space was reduced to one, where a patient sitting in a wheel chair was not recognized by the *Person detection* skill.

During these trials, the robot came in close contact with a total of 353 people (16.8 on average per training). Manual intervention for re-identification was performed 56 times. Again, most cases were sent stops due to uncertainty of the re-identification module. The issues identified are problems in person detection due to larger distances between robot and patient, and changing lighting conditions from hallway to darker corridors. Only in nine cases manual intervention via control tablet was necessary due to false decisions of the *Person re-identification* skill. Relating to the number of people standing nearby in these situations, this is about the same performance as achieved during the functional tests with staff members described in Sect. 7.1.3. The main reasons for this were very closely standing persons with similar appearance and false-positive person detections of the detector modules. To address this remaining re-identification issue, for future user trials we plan to additionally use an external technical device to be carried by the patient which is emitting ultrasound signals that can be located by the robot by analyzing the interaural time difference (ITD) of the arriving sound signals. To this end, the robot still has to be equipped with two additional ultrasound receivers. Fusing the decisions of both cues, the vision-based re-identification and the technical user tracking, should minimize manual interventions drastically and, additionally, enable a retraining of the appearance-based model of the current user.

Apart from these issues still to be solved, the robot coach has functioned technically robust in the clinical setting, and the robot's services were usable by the patients. Nevertheless, handling the robot's GUI and its services obviously needs some practice, and some functions need to be improved till the next user trials scheduled for end of the year 2015. Most involved patients appreciated the robot's instrumental function as walking coach, however, they rated the usefulness of the robot in its current state as limited due to the restricted range of training tours and the still rudimentary walking exercising.

For the last phase of the ROREAS project, which is still running till March 2016, a number of further user trials with stroke patients with increasing complexity regarding the offered training routes and the available human-robot interaction capabilities is already planned.

In these studies, the *usability* of the training application by the patients will be evaluated by our project partner from social sciences (SIBIS Institutes Berlin) using the indicators effectiveness, robustness, safety and joy of use. As it is difficult to reproduce exactly the same conditions for each run, a qualitative research design has been chosen. Data is going to be collected by participative observation, and participants (experts observing the patients' self-training with the robot) are asked to think aloud during the exercises. Subsequent to user trials, semi-structured interviews will be conducted with the patients and medical experts. Then we are going to see how well the robot's behaviors and offered training services fit into the self-training concept and can foster the physical and mental wellbeing of the stroke patients.

8 Summary and outlook

This paper describes the current state in developing and evaluating an assistive mobile robot coach for walking self-training of patients in clinical post-stroke rehabilitation. It gives an overview of the desired training scenario and the challenges and requirements of the clinical environment and presents the mobile robot coach ROREAS explicitly developed for this kind of robot-assisted walking training. Moreover, it describes the robot's functional system architecture and those HRI and navigation skills required for a robot coach that can operate autonomously in a clinical everyday environment. Then it reports on our three-stage approach in conducting function and user tests in the operational environment of the rehab center under everyday conditions. For security reasons, in the first stage we did not interact with real patients but replaced them by briefed technical staff of our robotics lab. Based on this, in the second stage we tested the walking coach with clinical staff imitating the walking behavior of stroke patients. As the walking coach had proven to be reliable and suitable, we started user trails with actual stroke patients at stage three. In these trials, the robot could fulfill almost all expectations regarding its navigation and interaction capabilities, nevertheless, there are diverse skills that need to be further advanced by algorithmic improvements and technical solutions to guarantee an autonomous walking training without the assistance of external observers making the control tablet obsolete. For the last phase of the project running till March 2016, a number of further user trials with stroke patients with increasing complexity regarding the offered training routes and the available human-robot interaction capabilities is planned.

The question, how such a robot assistant may practically be integrated in the clinical setting in the future is still an open issue, that depends on numerous factors, such as the usability and acceptance of the robot coach by the patients and the clinic staff, the therapeutic benefit of a robot-assisted self-training, the everyday and long-term suitability of the robot coach, its costs for acquisition and operation, new opportunities for the clearing of the training with health or pension insurance funds, safety regulations of the German Technical Inspection Agency (TÜV), or the certification of the robot coach as medical device. The ROREAS project can only try to clarify a few of these factors, and much work still remains to be done to transfer our project results into clinical practice. We are also aware, that any claims of real benefits of robotic assistance can only be substantiated by controlled comparative studies directly comparing robot-based assistive services to relevant conventional approaches (Andrade et al. 2014). The ROREAS project hopes to make a significant contribution by gathering information about the performance of assistive technology in real life and in daily clinical routine.

Acknowledgements

The authors wish to thank the partners of the ROREAS research consortium for their trustful cooperation allowing them to do robotics and HRI research in a challenging real-world scenario and environment, the "m&i Fachklinik" Rehabilitation Center Bad Liebenstein, the SIBIS Institute for Social Research in Berlin, and the health insurance fund Barmer GEK Wuppertal.

Compliance with Ethics Requirements

This work has received funding from the German Federal Ministry of Education and Research to the project ROREAS (grant agreement no. 16SV6133).

All authors declare that they have no conflict of interest.

All procedures followed were in accordance with the ethical standards of the responsible committee on human experimentation (institutional and national) and with the Helsinki Declaration of 1975, as revised in 2008 (5). Informed consent was obtained from all patients for being included in the study.

Additional informed consent was obtained from all patients for which identifying information is included in this article.

References

Andrade, A., Pereira, A., Walter, S., Almeida, R., Loureiro, R., Compagna, D. & Kyberd, P. (2014). Bridging the gap between robotic technology and health care. *Biomedical Signal Processing and Control*, 10(10), 65–78

Arras, K., Mozos, O. & Burgard, W. (2007). Using boosted features for the detection of people in 2D range data. In *IEEE Internat. Conf. on Robotics and Automation (ICRA)*, pp. 3402–3407

Balakrishnan, N. & Basu, A.P. (1996). Exponential Distribution: Theory, Methods and Applications. Gordon and Breach

Basteris, A., Nijenhuis, S., Stienen, A., Buurke, J., Prange, G. & Amirabdollahian, F. (2014). Training modalities in robot-mediated upper limb rehabilitation in stroke: a framework for classification based on a systematic review. *Journal of NeuroEngineering and Rehabilitation*, 11 (111), doi:10.1186/1743-0003-11-111

Einhorn, E., Langner, T., Stricker, R., Martin, C. & Gross, H.-M. (2012). MIRA - Middleware for robotic applications. In *EEE/RSJ Internat. Conf. on Intelligent Robots and Systems (IROS)*, pp. 2591–2598

Einhorn, E. & Gross, H.-M. (2015). Generic NDT mapping in dynamic environments and its application for lifelong SLAM. *Robotics and Autonomous Systems*, 69: 28-39

Eisenbach, M., Vorndran, A., Sorge, S. & Gross, H.-M. (2015). User recognition for guiding and following people with a mobile robot in a clinical environment. In *IEEE/RSJ Internat. Conf. on Intelligent Robots and Systems (IROS)*, pp. 3600–3607

Eisenbach, M., Kolarow, A., Vorndran, A., Niebling, J. & Gross, H.-M. (2015). Evaluation of Multi Feature Fusion at Score-Level for Appearance-based Person Re-Identification. In *Intenational Joint Conf. on Neural Networks (IJCNN 2015)*, pp. 469-476

EU-Robotics (2015). Robotics 2020 multi-annual roadmap for robotics in Europe - Horizon 2020. http://www.eu-robotics.net

Feil-Seifer, D. & Mataric, M. (2011). Socially Assistive Robotics. *IEEE Robotics and Automation Magazine*, 1:24–31

Fischinger, D., Einramhof, P., Papoutsakis, K., Wohlkinger, W., Mayer, P., Panek, P., Hofmann, S., Körtner, T., Weiss, A., Argyros, A. & Vincze, M. (2014). Hobbit, a care robot supporting independent living at home: First prototype and lessons learned. *Robotics and Autonomous Systems*, doi:10.1016/j.robot.2014.09.029

Fox, D., Burgard, W. & Thrun, S. (1997). The dynamic window approach to collision avoidance. *IEEE Robotics and Automation Magazine*, 4(1), 23–33

Gross, H.-M., Boehme, H.-J., Schroeter, C., Mueller, S., Koenig, A., Martin, C., Merten, M. & Bley, A. (2008). Shopbot: Progress in developing an interactive mobile shopping assistant for everyday use. In *IEEE Internat. Conf. on Systems, Man, and Cybernetics (SMC)*, pp. 3471–3478

Gross, H.-M., Boehme, H.-J., Schroeter, C., Mueller, S., Koenig, A., Einhorn, E., Martin, C., Merten, M. & Bley, A. (2009). TOOMAS: Interactive shopping guide robots in everyday use - final implementation and experiences from long-term field trials. In *IEEE/RSJ Internat. Conf. on Intelligent Robots and Systems (IROS)*, pp. 2005–2012

Gross, H.-M., Schroeter, C., Mueller, S., Volkhardt, M., Einhorn, E., Bley, A., Langner, T., Merten, M., Huijnen, C., van den Heuvel, H. & van Berlo, A. (2012). Further progress towards a home robot companion for people with mild cognitive impairment. In *IEEE Internat. Conf. on Systems, Man, and Cybernetics (SMC)*, pp. 637–644

Gross, H.-M., Debes, K., Einhorn, E., Mueller, S., Scheidig, A., Weinrich, C., Bley, A. & Martin, C. (2014). Mobile robotic rehabilitation assistant for walking and orientation training of stroke patients: A report on work in progress. In *IEEE Internat. Conf. on Systems, Man, and Cybernetics (SMC)*, pp. 1880–1887

Gross, H.-M., Mueller, S., Schroeter, Ch., Volkhardt, M., Scheidig, A., Debes, K., Richter, K. & Doering, N. (2015). Robot companion for domestic health assistance: Implementation, test and case study under everyday conditions in private apartments. In *IEEE/RSJ Internat. Conf. on Intelligent Robots and Systems (IROS)*, pp. 5992–5999

Hirata, Y., Hara, A. & Kosuge, K. (2007). Motion Control of Passive Intelligent Walker Using Servo Brakes. In *IEEE Transactions on Robotics*, 23(5), 981-990

Jacobs, T. & Graf, B. (2012). Practical evaluation of service robots for support and routine tasks in an elderly care facility. In *IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO)*, pp. 46–49

Jensen, B., Tomatis, N., Mayor, L., Drygajlo.& A, Siegwart, R. (2005). Robots meet humans - interaction in public spaces. *IEEE Trans. Industrial Electronics*, 52(6), 1530–1546

Kennedy, E.R. (1995). Particle swarm optimization. In *IEEE Internat. Conf. on Neural Networks (ICNN)*, pp. 1942–1948

Kulyukin, V., Gharpure, C. & Nicholson, J. (2005). RoboCart: Toward Robot-Assisted Navigation of

Grocery Stores by the Visually Impaired. In *IEEE/RSJ Internat. Conf. on Intelligent Robots and Systems (IROS)*, pp. 2845–2850

Magnusson, M., Nüchter, A., Lörken, C., Lilienthal, A. & Hertzberg, J. (2009). Evaluation of 3D registration reliability and speed - a comparison of ICP and NDT. In *IEEE Internat. Conf. on Robotics and Automation (ICRA)*, pp. 3907–3912

Quigley, M., Gerkey, B., Conley, K., Faust, J., Foote, T., Leibs, J., Berger, E., Wheeler, R. & Ng, A. (2009). ROS: an open-source robot operating system. In *ICRA Workshop on Open Source Software*

Philippsen, R. & Siegwart, R. (2005). An interpolated dynamic navigation function. In *IEEE Internat. Conf. on Robotics and Automation (ICRA)*, pp. 3782–3789

Rodriguez-Losada, D., Matia, F., Jimenez, A., Galan, R. & Lacey, G. (2005). Implementing map based navigation in Guido, the robotic smartwalker. In *IEEE Internat. Conf. on Robotics and Automation (ICRA)*, pp. 3390–3395

Ross, A. & Nandakumar, K. (2009). *Encyclopedia of Biometrics*. pp. 611–616, Springer

Stoyanov, T., Magnusson, M., Almqvist, H. & Lilienthal, A. (2011). On the Accuracy of the 3D Normal Distributions Transform as a Tool for Spatial Representation. In *IEEE Internat. Conf. on Robotics and Automation (ICRA)*, pp. 4080–4085

Schroeter, Ch., Mueller, S., Volkhardt, M., Einhorn, E., Huijnen, C., van den Heuvel, H., van Berlo, A., Bley, A. & Gross, H.-M. (2013). Realization and user evaluation of a companion robot for people with mild cognitive impairments. In *IEEE Internat. Conf. on Robotics and Automation (ICRA)*, pp. 1145–1151

Shiomi, M., Zanlungo, F., Hayashi, K. & Kanda, T. (2014). Towards a Socially Acceptable Collision Avoidance for a Mobile Robot Navigating Among Pedestrians Using a Pedestrian Model. In *International Journal on Social Robotics*, 6(3), 443-455

Stricker, R., Mueller, S., Einhorn, E., Schroeter, C., Volkhardt, M., Debes, K., Gross, H.-M. (2012). Konrad and Suse, two robots guiding visitors in a university building. In *Autonomous Mobile Systems (AMS)*, Informatik Aktuell, pp. 49–58, Springer

Triebel, R., Arras, K., Alami, R., et al. (2015). SPENCER: A Socially Aware Service Robot for Passenger Guidance and Help in Busy Airports. In *Field and Service Robotics (FSR)*

Trinh, T., Schroeter, Ch., Kessler, J., Gross & H.-M. (2015). "Go ahead, please": Recognition and resolution of conflict situations in narrow passages for polite mobile robot navigation. In *International Conf. on Social Robotics (ICSR)*, (pp.643-653)

Volkhardt, M., Weinrich, C. & Gross, H.-M. (2013). People tracking on a mobile companion robot. In *IEEE Internat. Conf. on Systems, Man, and Cybernetics (SMC)*, pp. 4354–4359

Wade, E., Parnandi, A., Mead, R. & Mataric, M. (2011). Socially assistive robotics for guiding motor task practice. *Journal of Behavioral Robotics*, pp. 218–227

Weinrich, C., Vollmer, C. & Gross, H.-M. (2012). Estimation of human upper body orientation for mobile robotics using an svm decision tree on monocular images. In *IEEE/RSJ Internat. Conf. on Intelligent Robots and Systems (IROS)*, pp. 2147–2152

Weinrich, C., Wengefeld, T., Schroeter, C. & Gross, H.-M. (2014). People detection and distinction of their walking aids in 2D laser range data based on generic distance-invariant features. In *IEEE Internat. Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 767–773

Weiss, A., Bernhaupt, R. & Tschegili, M. (2011). The USUS Evaluation Framework for user-centred HRI. *New Frontiers in Human-Robot-Interaction*, pp. 125-147, John Benjamins

Horst-Michael Gross is full professor for Computer Science at Technische Universität Ilmenau and head of the research lab for Neuroinformatics and Cognitive Robotics. He received his Diploma degree (M.Sc.) in Technical & Biomedical Cybernetics (1985) and his Doctorate degree (Ph.D.) in Neuroinformatics (1989) from the Ilmenau University. His research mainly focuses on the development of robust and adaptive techniques for mobile navigation and HRI that cover aspects such as mapping, SLAM, multi-modal user detection and tracking, human-aware navigation, situation recognition, and dialog adaptation. As challenging application field, his research is focused on socially assistive mobile robots for public and domestic environments. Research highlights of his lab are the shopping guide robot TOOMAS (2000-2008), a robot companion for patients suffering from MCI as part of the EU-project CompanionAble (2008-12), the mobile health robots Max and Tweety developed within the SERROGA project (2012-15), and the robot coach ROREAS for walking training of stroke patients (2013-16).

Andrea Scheidig is postdoctoral assistant at the research lab for Neuroinformatics and Cognitive Robotics of the Technische Universität Ilmenau. She received her Diploma degree (M.Sc) in Computer Science (1996) and her Doctorate degree (Ph.D.) in Neuroinformatics (2003) from the Ilmenau University. Her main research interests are socially assistive robots in clinical and domestic application scenarios, benchmarking of assistive robots in real environments, and user-adaptive human-robot interaction.

Klaus Debes is a senior lecturer at the research lab for Neuroinformatics and Cognitive Robotics of the Technische Universität Ilmenau. He received his Diploma degree (M.Sc.) in Technical and Biomedical Cybernetics (1978) and his Doctorate degree (Ph.D.) in Biomedical Engineering (1984) from the Ilmenau University. In 1988, he joined the Xi'an Jiaotong University in China as a postdoctoral assistant for one year. His main research interests are soft computing and machine learning for industrial applications and cognitive robotics.

Erik Einhorn is Ph.D. student at the research lab for Neuroinformatics and Cognitive Robotics at the Technische Universität Ilmenau. He received his Diploma (M.Sc.) degree in Computer Science from this university in 2007. His research interests are focussed on visual 3D perception, environment modeling, and mapping for mobile robot navigation.

Markus Eisenbach received his Diploma degree (M.Sc.) in Computer Science from Technische Universität Ilmenau in 2009. After this, he joined the Neuroinformatics and Cognitive Robotics Lab at this university as research assistant. Currently he is doing a Ph.D. in robotics. His research interests are robotics, machine learning, computer vision and appearance-based person re-identification in particular, and deep learning.

Steffen Müller received his Diploma degree (M.Sc.) in Computer Science from the Technische Universität Ilmenau in 2005. Then he joined the Neuroinformatics and Cognitive Robotics Lab at this university as research assistant. Currently he is doing a Ph.D. in robotics. His research interests are mobile robotics, personalization of Human-Robot Interaction, and adaptation of robot behavior by learning.

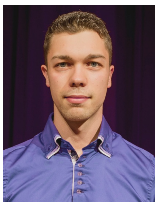

Thomas Schmiedel received his B.Sc. in Information and Communication Technology from the HFT Leipzig University of Applied Sciences (Germany) in 2011. Then, he did a Master in Computer Science at the Technische Universität Ilmenau and graduated in 2014. After this, he joined the Neuroinformatics and Cognitive Robotics Lab at this university as a research assistant, where he is currently working on his Ph.D. thesis. His research interests include robot navigation, depth perception, and machine learning.

Thanh Q. Trinh holds a Bachelor's and Master's degree in Computer Science from Technische Universität Ilmenau, which he received in 2011 and 2014, repectively. After this, he joined the Neuroinformatics and Cognitive Robotics Lab at this university as a research assistant, where he is currently working on his Ph.D. thesis. His research interests primarily focus on situation- and human-aware navigation for a polite human-robot interaction.

Christoph Weinrich received a Diploma Degree (M.Sc.) in Computer Science from the Technische Universität Ilmenau in 2009. Then he joined the Neuroinformatics and Cognitive Robotics Lab of this university as PhD-student and worked on research topics in the field of laser-based people detection and vision-based pose estimation as well as socially compliant navigation of assistance robots. He received his Ph.D. in Robotics from the Technische Universität Ilmenau in 2016.

Tim Wengefeld received his Bachelor's and Master's Degree in Computer Science from Technische Universität Ilmenau in 2010 and 2014, respectively. He is currently employed at the Neuroinformatics and Cognitive Robotics Lab, while doing his Ph.D. in robotics. His research is focused on people-detection and -tracking, machine learning and pattern recognition.

Andreas Bley is one of the founders and CEO of the company MetraLabs GmbH. He studied in Industrial Engineering and Management with a focus on automation, intellectual property rights, and finance at Technische Universität Ilmenau between 1998 and 2003. In 2007, he completed his doctoral thesis about the usage of renewable energies by private households at the Friedrich-Alexander-University Erlangen-Nürnberg. During his studies he held a scholarship of the Klaus Murmann Fellowship Programme of the Foundation of German Business. Fascinated by interaction of service robots and humans, he managed research projects in this topic at MetraLabs, e.g. EU FP7 CompanionAble or EU FP7 Robot-Era.

Christian Martin is one of the founders of the company MetraLabs GmbH, where he is CEO and Head of Software Development. He received a Diploma Degree (M.Sc.) in Computer Science from the Technische Universität Ilmenau in 2004. From 2004 to 2007, he worked as research assistant at the Neuroinformatics and Cognitive Robotics lab at this university. He received his Doctoral degree (Ph.D.) in Computer Science and Cognitive Robotics from Ilmenau University in 2012. His current research interests are Cognitive Robotics (autonomous robots, human robot interaction, probabilistic robotics) and control architectures for autonomous mobile robots.