in: Proc. 21st IEEE Int. Symposium on Robot and Human Interactive Communication (Ro-Man 2012), Paris, France, pp. 695-700, IEEE 2012

Interactive Mobile Robots Guiding Visitors in a University Building

Ronny Stricker, Steffen Müller, Erik Einhorn, Christof Schröter, Michael Volkhardt, Klaus Debes, Horst-Michael Gross

***Abstract*— This paper presents an architectural overview of a robot-based visitor information system in a university building. Two mobile robots serve as mobile information terminals providing information about the employees, labs, meeting rooms, and offices in the building and are able to guide the visitors to these points of interest. The paper focuses on the different software components needed to meet the requirements of the multi-story office building. Furthermore, the integration of a multi-hypotheses person tracker is outlined, which helps the robots to interact with the people in their near surrounding. Besides first observations on interaction, the further development is outlined as well.**

I. Introduction

The long-term research project Zuse-Guide aims at developing two autonomous interactive service robots called “Konrad” and “Suse” which are to autonomously guide and to tour visitors within the new Konrad Zuse building of the Ilmenau University of Technology, hosting the School of Computer Science and Automation on 4 floors since June 2011. The service robots are supposed to provide an intuitive interaction to potential users. Service tasks considered in this project are to autonomously navigate in the building and guide the visitors, to provide information about the different labs and to tour visitors around in the building, giving them a deeper insight into the research topics and exhibits of the institutes and labs. During their operation, the robots wait next to the main entrance doors at the ground floor and offer their service when a potential user approaches. The user can select the different functions using a touch screen-based graphical dialog system which is commented by a speech synthesis system to provide a natural and intuitive user experience. The robots also offer to guide the visitor to the destination of his/her choice and inform about news for the different institutes and labs. When the robots have finished guiding the user around, they return to their waiting position to welcome the next visitor.

The main part of the visitor information system are the two guiding robots (Fig. 1). The robot platform is a SCITOS A5 produced by Metralabs GmbH Ilmenau, which originally has been developed for the application as shopping assistant [1].

Besides the main sensors and devices used in the shopping robot setup, the robot is equipped with an RFID antenna at the bottom, able to detect tags burrowed in the floor (for emergency stop in front of the stairways). Additionally, for operating the elevator, an infrared sender has been mounted that allows remote calling of the elevator and the “manual selection” of the target story. Although the features offered at the current state of implementation are still limited, especially in comparison to our previous work on a mobile shopping guide robot presented in [2], in this project we are trying to focus on a robust and fault tolerant system, running continuously during working hours five days a week. In the near future, the robots will undergo iterative extension and will serve as test platforms for new methods in the fields of navigation and human-robot interaction in public environments.

All authors are with the Neuroinformatics and Cognitive Robotics Lab, Ilmenau University of Technology, 98693 Ilmenau, Germany http://www.tu-ilmenau.de/neurob

This paper summarizes the project objectives as well as the current state of implementation which already allows for a continuous guidance service. Since this project is still ongoing work, the paper concentrates on the aspect of the robot infrastructure and the functional interaction of the different software modules. Therefore, we will give an overview of the constraints and requirements arising from this scenario in the following Section. After that, we will give an insight into the infrastructure of the building in Sec. III. The administrative software modules and a detailed overview of the different software modules running on the robot platform is given in Sec. IV with a focus on the multi-story navigation planner and the person tracker. The paper concludes with some findings on the first days of operation and an outlook.

II. Challenges and Constraints of the Application Scenario and State-of-Art

In the last years, diverse mobile service robots employed as tour guides for exhibitions and museums have been presented [3]. Among them are such well known robots as

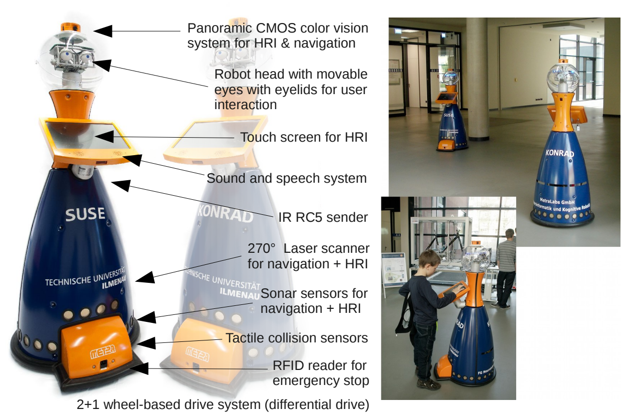

Fig. 1. Left: sensors and interaction devices of the robot platforms, Right: Konrad and Suse waiting at the main entrance and Konrad during open house presentation.

in: Proc. 21st IEEE Int. Symposium on Robot and Human Interactive Communication (Ro-Man 2012), Paris, France, pp. 695-700, IEEE 2012

Rhino, Minerva, and Sage, the exposition guide RoboX, or the robots Mona/Oskar at the Opel sales center in Berlin. Also Fraunhofer IPA developed robots for the museum for communication in Berlin [4]. Usually, these robots guide visitors to a set of exhibits while offering related information, and thus have a similarity to our application. However, none of those systems is working in a multi-story environment and is involved in such a complex infrastructure, like a School of Computer Science with hundreds of employees, umpteen of offices and labs, and a highly dynamic building layout.

The Konrad Zuse building has more than 200 rooms (labs, lecture rooms, offices) and accommodates more than 225 people. In total, the building has an area of more than 11,000 m^2 on 4 floors, of which the hallways the robots operate in are about 5,000 m^2 (Fig. 2).

More challenges arise from the multi-story architecture of the building. Not only must the robots be able to navigate over different floor levels, but even more important, the stairways are dangerous obstacles which have to be avoided very robustly. Multi-story navigation with incorporation of elevator usage is known from mobile transport systems [5], [6], which usually are centrally controlled or using fixed pathways for navigation. Since our robots are supposed to navigate fully autonomously reacting on the users' requirements, they have to be able to control and use the elevators of the building, and to recognize the current floor they are operating on. That means multi-story localization is required.

Our operational environment is highly dynamic. Staff working in the labs and students are moving and working in the public areas. There are changing furniture and exhibits, and experimental setups are occupying the hallways, resulting in a quite crowded environment at some times. All this is requiring a robust and situation-aware navigation.

For unsupervised continuous operation, docking stations enable the robots to autonomously recharge their batteries. Such self-sustaining operation can also be found at the shopping assistants of our former projects [2], [7].

Since the robots operate in a large area were it might take some time to find them when they are not at their initial waiting position, a convenient monitoring and remote control interface is required - both for the system administrators and the potential users.

Besides these technical aspects, the human-machine interaction also plays an important role in our scenario. People have to be made interested in using the information services, and communication should be adaptive and situation aware. Various studies regarding this topic, again have been conducted in the domain of entertainment and tour guide robots, as for example in [8] or in [9] which dealt with the classification of potential user according their interest or need for help based on the movement trajectories.

III. Operational Area

As already described in the introduction, the public area of a large multi-story office building comes along with numerous challenges. This section will introduce the environmental conditions and necessary installations to enable a safe operation of autonomous robots.

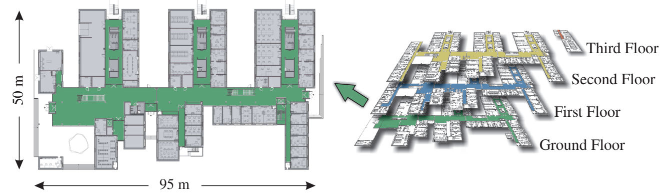

Fig. 2. Different floors of the four-story Zuse building with a detailed view of the ground floor (operation area for the guide robots highlighted in color). The first and second floor have a similar shape, whereas the third floor only comprises two lecture rooms.

Resulting from the four floor architecture, the central aisle shows a couple of stairs, that are the most critical areas for the robots. Downward stairs pose the risk of fall, while passing below an upward stair might result in a collision with the robot's head. The limited vertical range of the laser range scanner and sonar sensors does not allow to perceive both types of obstacles. In order to prevent crashes even when a faulty localization leads the robot to the proximity of these obstacles, the robot is equipped with an RFID reader, closely coupled to the drive system. RFID tags burrowed in the floor at the critical areas (in front of and beneath stairs) can deactivate the drive in less than one second. Tags have been placed in three rows to guarantee at least one detection when passing over them even at high speed. In addition to the RFID based stair detection, we are currently preparing a second and independent system based on the detection of magnetic tapes which are burrowed in front of the stairs to meet the safety requirements of public buildings and the criteria of the German Technical Inspection Agency (TÜV). As long as the installation of these is not finished, for safety reasons we are limited to operate on the ground floor only at the moment. As a consequence of this temporary limitation, the user can only be guided to the stairs or to one of the elevators if the target is located at one of the upper floors.

A second requirement is the ability to operate elevators. In order to enable the robots to control the elevators with only minimal changes to the control electronics of the elevators themselves, each of the 7 elevator doors was equipped with an RC-5 receiver (a protocol for consumer infrared remote control) that is connected to the call button of the existing elevator control electronics. In the same manner, two RC-5 receivers are installed in each of the two elevator cabins connected to the buttons for selecting the destination floor. In effect, robots interact with the elevator nearly similar to humans, and no change in the existing elevator controllers was necessary which is not allowed due to warranty terms.

IV. Software Architecture

A. Administrative Backend

Remote Administration and the surveillance is a very important issue when running robots fully autonomously. Therefore, the robots are integrated in a framework which enables the communication with the robots. This section gives a brief introduction into the used communication

in: Proc. 21st IEEE Int. Symposium on Robot and Human Interactive Communication (Ro-Man 2012), Paris, France, pp. 695-700, IEEE 2012

framework and explains the different software components involved to build up the background for the information system.

As shown in Fig. 3, the central administrative component is a server hosting a MySQL database and a webserver. Furthermore, on that host a robot server is running for each of the robots, which is the server side communication endpoint for the robot. Software on the robots as well as the robot server are based on the MIRA middleware (www.mira-project.org), providing inherent means for sharing robot data between distributed modules. Although, a number of robot middlewares have been developed e.g. ROS or YARP, most of them share the same disadvantage of a centralized approach [10]. Since the robots should act as independent systems, which should only be observed by a centralized server, we don't want to introduce this single point of failure. Furthermore, MIRA has some advantages concerning RPC functionality, transmission of complex datatypes and CPU usage [10].

The administration server has two main functions. On the one hand, it offers a website interface for maintaining the knowledge base comprising room data, institute and lab information, exhibits, and people working in the building. Authorized staff can edit department specific news pages and change assignments of people to rooms and exhibits. In the future, it will also be possible to book the robot for tours or welcoming guests at the main entrance at a certain time via this website. All this data is stored in one database, that is accessed by the robots each time a user requests the respective information. On the other hand, there is an administration website for remotely observing and controlling the state of the robots. A floor plan indicates the current position, connection state and internal state of the robots, and allows to send the robots to arbitrary target positions.

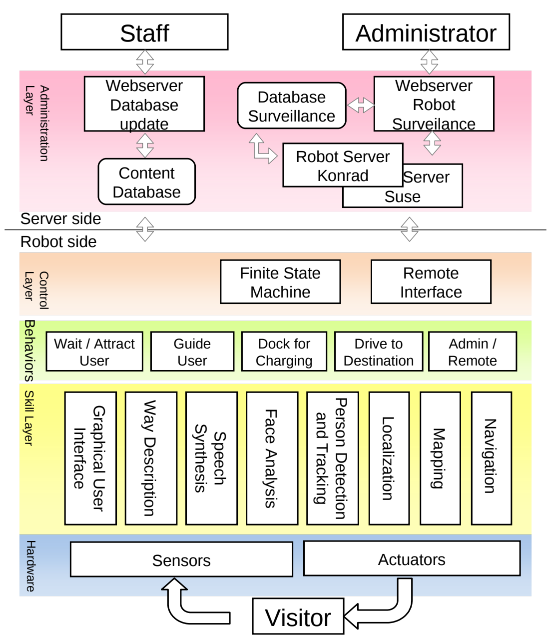

Fig. 3. System architecture, consisting of mobile robots part (bottom) and server side (top)

B. *Robot-side software modules*

Relying on the administrative infrastructure, each robot operates an application that is also based on the MIRA middleware framework with the aim to guarantee the main requirements to a modern robot control architecture, like modularity, extensibility, efficiency, customizability, reusability, and rapid application development. Similar to the former software architecture used for the shopping robots [1], the robotics related methods and skills have been abstracted from the application itself resulting in a flexible layered system architecture (Fig. 3). In this architecture, the bottom-most Hardware Layer encloses the hardware (sensors and actuators) and the operating system. The low-level sensor information is processed in the next level, the Skill Layer, which provides a set of necessary, robotic-specific basic skills. These skills are only weakly coupled and do not need any coordination since they are operating in parallel using exclusive resources. The Skill Layer covers the whole spectrum of navigation skills including map building and localization, as well as person detection and tracking. It is planned to also integrate face analysis [11] in order to get information on the interest, attention or moods of the current user and enable an adaptation of the robot's interaction behavior. Speech synthesis and the GUI, as well as the path description generation are also placed in this layer and provide their services to the higher application levels. Above the independent skills, there is a finite state machine representing the application control which makes use of the basic features provided by the skills and is orchestrating the behaviors. These behaviors are exclusive units each representing an individual control loop for accomplishing the different task-oriented navigation and interaction functions of the robot.

1) Localization and Navigation: Since our operational area has four different floors, localization and path planning can not be performed using a single occupancy grid map. Instead, we use a hybrid, hierarchical topological map for path planning and localization. The topological map consists of nodes and gateways that are connected by directed edges (see Fig. 4).

Each node represents a certain region within the world, e.g. rooms or floors. These nodes build a hierarchy, i.e. nodes can be subdivided into sub-nodes or grouped into parent-nodes. This allows us to model our building using different levels of detail. On the coarsest level, each node represents a single floor of the building. Each such node is further subdivided into sub-nodes that correspond to the aisles of each floor. These nodes are further subdivided into nodes that represent rooms, corridors, etc. On the finest level, the leaf nodes contain local metric occupancy maps of the rooms.

The gateways represent transitions between different regions. Typical examples for gateways are doors or elevators. Unlike nodes, gateways have a metric position. Their position usually is given within the metric occupancy grid maps stored in the leaf nodes. Gateways therefore are the links between the topological and the metrical information in the hybrid

in: Proc. 21st IEEE Int. Symposium on Robot and Human Interactive Communication (Ro-Man 2012), Paris, France, pp. 695-700, IEEE 2012

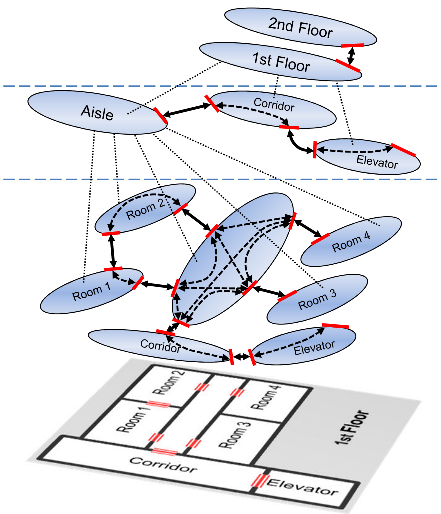

Fig. 4. The different levels of a hybrid topological map. Nodes (blue) represent regions of the environment and can be subdivided into subnodes. Each node has associated gateways (red) that are connected to other gateways. Gateways of the same node are fully connected (dashed arrows). Connections between gateways of different nodes represent transitions between different topological nodes (solid arrows). On the finest level of hierarchical map, the nodes are associated to occupancy grid maps.

map. They define the metric locations, where the robot can move from one topological node to another.

Gateways are connected to each other via directed edges. The edges are associated with traversal costs that are taken into account for path planning. Here we can differentiate between two different edge types:

1. Gateways within the same node are fully connected to each other via bidirectional edges. In other words, we assume that the robot can freely move between the different doors (the gateways) of a single room (the node). The traversal costs between gateways within leaf nodes are computed in a preprocessing step directly on the underlying metric occupancy maps using the E* planning algorithm [12]. These traversal costs are propagated upwards to the parent nodes in order to compute the traversal costs between the gateways on a coarser level.

2. Gateways of different nodes are connected using directed edges. We call these edges transitions since they allow the robot to move from one topological node to another. The costs for these transitions can be chosen arbitrarily and are usually set to zero, since driving through a door does not impose any costs. Nevertheless, using these costs, one can model specific properties of the building. In a building which, for example, has a fast and a slow elevator, the slow elevator should have higher transition costs. In order to minimize the time that is necessary to reach the goal, the robot may then use the fast elevator even though the path to the goal may become longer compared to using the slower elevator.

When planning a path to a goal, we search the shortest path between the leaf-node, the robot is currently located in, and the leaf-node that contains the goal. Therefore, the planning starts on the coarsest level of the topological map using a general Dijkstra algorithm that takes the costs between the gateways into account. The planning process is iteratively refined using the sub-nodes on levels with higher detail until the two leaf-nodes of the robot's current location and the goal position are reached. Here, the planning continues on the stored occupancy grid map using the computationally more expensive E* path planning algorithm [12]. This hierarchical approach combines fast path planning on a topological map and the ability of dynamic replanning that is supported by the E* algorithm. The result of the planning process is a list of gateways, the robot needs to traverse in order to reach the goal.

Nodes and gateways can also have certain semantic actions associated, that are sent to the application layer when the gateway is traversed or a node is entered by the robot. Such actions include interactions with the building, like calling the elevator, waiting for the elevator door to open, giving context information to the user, etc. This allows us to integrate such high-level actions seamlessly into the robot navigation.

For localization, we use a particle filter-based Monte Carlo localization, that operates on the metric occupancy maps stored within the leaf nodes of the hybrid topological map. As sensors, we use the laser range finder (see Fig. 1) for localization within those occupancy maps. For the estimation of the current floor, the robot uses an embedded accelerometer, that allows to estimate the vertical movement of the robot. To get rid of drifting issues introduced by the integration of the vertical movements, the sensor values are evaluated if the robot stands within the elevator cabin only.

Due to the variety of the different behaviors, simple navigation tasks like driving to a given position are not sufficient. Instead, additional tasks like searching and following a person, or autonomous docking to a docking station must be supported. These different tasks must be handled without reconfiguration or restart, since all of these navigation behaviors are part of the user guiding application. To meet all of these requirements, we have developed a modular navigation architecture [13]. Such a modular navigator requires a suitable interface allowing to specify complex tasks. Therefore, we introduced a task-based system which allows us to define navigation jobs consisting of several sub-tasks and their corresponding parameters, such as the goal point to drive to, the preferred driving direction of the robot, the accuracy for reaching the goal, etc. [14]. Thus, a task describes a combination of abstract navigation goals that should be fulfilled by completing all subtasks. In order to generate motion commands suitable to the currently active task, a motion planner based on the Dynamic Window Approach (DWA) [15] chooses actions from the range of possible velocities, by evaluating so-called objectives. An objective is a quality measure for a potential action, that is enabled or disabled based on the active subtasks and configured by their parameters. Therefore, objectives are concrete implementations of action evaluation functions specialized for certain abstract subtasks. This separation of abstract (sub)tasks and concrete objectives has shown to be advantageous for the navigator's flexibility and extensibility. The DWA planner combines the outputs of all active objectives in order to find the best action in the current situation and sends it to the motor controllers.

in: Proc. 21st IEEE Int. Symposium on Robot and Human Interactive Communication (Ro-Man 2012), Paris, France,
pp. 695-700, IEEE 2012

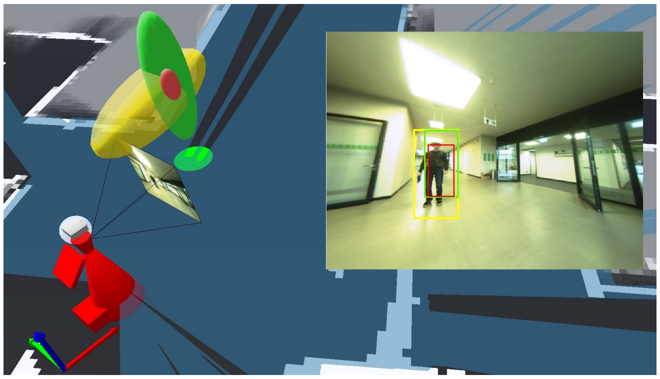

Fig. 5. MIRA visualization of exemplary detections and tracked hypothesis. 3D view of the robot, its camera image, its laser range scan (blue), an HOG detection with high variance in distance (yellow), a leg detection on the ground (green circle) and its aligned detection with high variance in height (green ellipsoid). The tracked hypothesis (red) includes all aggregated detections (motion and face not shown for clearity). The small image (2D view) shows the camera image with bounding boxes of the detections and projected hypothesis in respective coloring.

2) Person Tracking: The robot needs the ability to know the position of people in its close environment. This is useful for attracting people passing by [9] and for facing the detected persons with the pan tilt head (see Fig. 1) during interaction or a guiding tour. Furthermore, tracking and the subsequent prediction of person trajectories is necessary to enable socially acceptable navigation [16] which will be integrated in near future. The developed probabilistic multi-hypothesis people tracking system is based on a 6D Kalman Filter that tracks the position and velocity of people assuming an uncertain random acceleration [17]. The person tracker processes the detections of different, asynchronous observation modules – namely a laser-based leg detector, a face detector, a motion detector and two body shape detectors. The leg detector is based on the boosted classifier approach of [18] and the face detection system utilizes the well-known AdaBoost detector of Viola&Jones [19]. The motion detection cue is only active when the robot is standing still and utilizes a fast and simple image difference approach [20]. Finally, we apply a combination of a full body and an upper body shape detector based on Histograms of Oriented Gradients (HOG) [21], [22]. In case of visual detection modules, we transform the bounding box of the user into a 3D Gaussian detection by using the parameters of the calibrated camera. The distance of the 3D Gaussian is estimated by the size of the bounding box assuming a predefined metric width of the person detection. The sensor model is incorporated into the covariance of the Gaussian distribution resulting in a high variance in distance estimation for visual detections.

Each module detects persons by different body parts, i.e. the face, legs, or head-shoulder contour. Therefore, we use the head of people as a reference point for alignment. In this step, the positional variance on the vertical axis is increased according to the uncertainty of the head position to the detected body part, e.g. high additional variance for leg detections accounting for different heights of people but none for face detections (Fig. 5). Each detection is then transformed into a world coordinate frame using the transformation framework of MIRA middleware [10] which allows for a linear motion model in the Kalman Filter. This transformation also respects the uncertainty of the robots position given by the Monte Carlo localization. The resulting detections are associated with the closest predicted hypotheses in the system. If a distance threshold is exceeded, a new hypothesis is inserted at the position of the detection, while in the other case the Kalman Filter update is applied in order to improve the estimated position. The resulting set of person hypotheses and their history as the movement trajectory of the respective person are used by the higher level modules to make decisions on interaction activities.

3) Interaction and Control: Based on the person tracking and navigation skills, the central state machine can operate closely coupled to the graphical user interface (GUI). The state machine comprises 20 states where each state is associated with one of the behaviors (see Fig. 3). Transitions between states are triggered by navigation events, person tracking events, events from the GUI, or via the administration remote interface.

The behaviors are independent control loops, which are realizing either standing still and facing the detected persons with the pan-and-tilt head, or any driving activities. Examples are driving to a desired target with or without observing the presence of a person following, docking to the charging station and resting there, or driving remotely-controlled in the admin mode.

One central part of the search functionality is the generation of natural language path descriptions from arbitrary start points to a desired destination. This is based on a handcrafted labeled directed graph, that holds sentence parts on each edge. For a path description, the start and end nodes are selected based on the current position and orientation and a shortest path is generated through the graph minimizing the number of edges used. To find a description finally, the sentences on the path edges are concatenated by using a generic set of fill words and a random start phrase.

V. Results and Outlook

Our long-term trials have been started at the end of April 2012. The first three weeks of operation included an open house presentation during which a lot of experimental setups were occupying the hallways, and where large groups of people were passing by. Even under such circumstances, the robots were able to offer their services without problems. During the whole day, we have registered more than 150 interactions per robot with visual and verbal path description, leading to more than 60 successful guiding events per robot. The large gap between the number of interactions and the successful guiding events is caused by the fact that some users were satisfied when getting the plain path description and did not start a guidance tour. For a regular office day, when the robot typically operates for about 6 hours before it drives to its docking station, the numbers are lower. Therefore, we counted about 30 guiding tour events with 60 interactions per robot (Fig. 6).

in: Proc. 21st IEEE Int. Symposium on Robot and Human Interactive Communication (Ro-Man 2012), Paris, France, pp. 695-700, IEEE 2012

TABLE I

AVERAGE SCORES (0=WORST, 1=BEST) FOR THE EVALUATION OF SERVICE OFFERED BY THE ROBOT.

Question	Average value
Did you like my driving skills?	0.70
Could you use my touch screen intuitively?	0.83
Did you find my service useful?	0.61/0.84

After the service has been offered by the robot, the user is asked to answer three questions (Table I). To this end, the GUI provides a slider for every question, which can be moved freely between a sad looking smiley and a happy smiley. The selected slider position is interpreted as a value between 0 (worst) and 1 (best). Having a closer look at the results, we noticed that some users were very unhappy with the provided usefulness of the service. This can be easily explained, as due to the navigation restriction to the ground floor, these users were guided to the nearest elevator or stairs only (just a few meters away), since their goal was located on one of the upper floors. The users who were guided all the way to the target room were quite happy with the service, resulting in an average value of 0.84. Therefore, we expect these values to improve significantly as soon as the robot is allowed to drive to the upper floors.

Continuing our work, we are planning to extend the functionality of the robot step by step. The short-term objective is to enable the usage of all four floors and the elevators by integrating the magnetic tape detection to meet the safety requirements of public buildings. Furthermore, we will improve the continuous tracking of people during driving. However, the long-term objectives are to provide a person specific and adaptive guidance system, which can provide user-adaptive tours. The task arising from this objective comprise advanced user estimations, like gender, age, or interest in continuing the interaction, that allows for user group classification and a respective adaptation of interaction behaviors. Furthermore, the robot should be more entertaining by providing some kind of small talk and advertise exhibits if they are passed on the path to the destination.

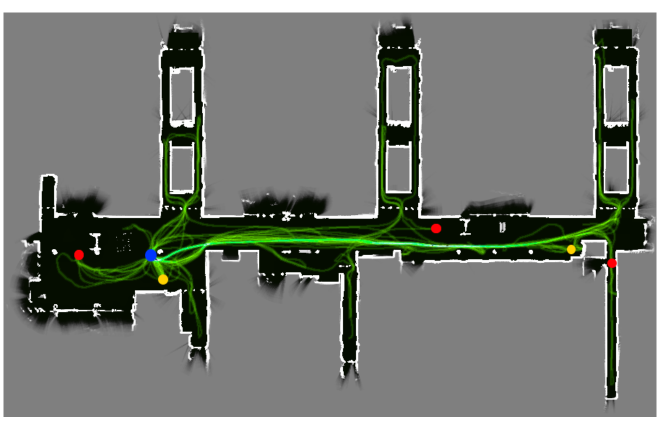

Fig. 6. Recorded track of one robot on a usual office day displayed on the occupancy map of the ground floor (hallways only). The waiting position of the robot in the entrance area is colored in blue. If the guide target is in one of the upper floors the user is guided to the stairs (colored red) or one of the elevators (highlighted in yellow). The tracks are different each time the robot drives to the target depending on the present dynamic obstacles.

REFERENCES

[1] H.-M. Gross, H.-J. Böhme, C. Schröter, S. Müller, A. König, Ch. Martin, M. Merten and A. Bley. ShopBot: Progress in Developing an Interactive Mobile Shopping Assistant for Everyday Use. In IEEE Int. Conf. on SMC, pp. 3471-3478, 2008.

[2] H.-M. Gross, H.-J. Böhme, C. Schröter, St. Müller, A. König, E. Einhorn, Ch. Martin, M. Merten and A. Bley. TOOMAS: Interactive Shopping Guide Robots in Everyday Use - Final Implementation and Experiences from Long-term Field Trials. In IEEE/RJS Intern. Conf. on Intelligent Robots and Systems, pp. 2005-2012, 2009.

[3] B. Jensen, N. Tomatis, L. Mayor, A. Drygajlo and R. Siegwart. Robots meet Humans - Interaction in public Spaces. In IEEE Trans. Industr. Electronics 52, pp. 1530-1546, 2005.

[4] R.D. Schraft, B. Graf, A. Traub and D. John. A mobile robot platform for assistance and entertainment. In Industrial Robot: An International Journal, vol. 28, no. 1, pp. 29-35, 2001

[5] K. Niechwiadowicz and Z. Khan. Robot Based Logistics System for Hospitals-Survey. In IDT Workshop on Interesting Results in Computer Science and Engineering, 2008.

[6] N. Tschichold-Gürman, SJ Vestli and G. Schweitzer. The service robot MOPS: First operating experiences. In Robotics and Autonomous Systems, vol. 34, no. 2, pp. 165-173, 2001.

[7] H.-M. Gross, Ch. Schröter, S. Müller, M. Volkhardt, E. Einhorn, A. Bley, T. Langner, M. Merten, C. Huijnen, H. van den Heuvel, A. van Berlo. Further Progress towards a Home Robot Companion for People with Mild Cognitive Impairment. In: Proc. IEEE Int. Conf. on Systems, Man, and Cybernetics (IEEE-SMC 2012), Seoul, IEEE 2012.

[8] F. Faber, M. Bennewitz, C. Eppner, A. Gorog, C. Gonsior, D. Joho,M. Schreiber and S. Behnke. The humanoid museum tour guide Robotinho. In The 18th IEEE International Symposium on Robot and Human Interactive Communication, pp. 891-896, 2009.

[9] S. Mueller, S. Hellbach, E. Schaffernicht, A. Ober, A. Scheidig, H.-M. Gross. Whom to talk to? Estimating user interest from movement trajectories. in: Proc. of the RO-MAN, pp. 532-538, 2008.

[10] E. Einhorn, T. Langner, R. Stricker, Ch. Martin, and H.-M. Gross. MIRA - Middleware for Robotic Applications. In IEEE/RSJ International Conference on Intelligent Robots and Systems, 2012.

[11] C. Martin and H.-M. Gross. A Real-time Facial Expression Recognition System based on Active Appearance Models using Gray Images and Edge Images. In 8th IEEE Int. Conf. on Face and Gesture Recognition, paper no. 299, 2008.

[12] R. Philippsen and R. Siegwart. An Interpolated Dynamic Navigation Function. In Proc. of the ICRA, pp. 3782-3789, 2005.

[13] E. Einhorn and T. Langner. Pilot - Modular Robot Navigation for Real-World-Applications. In Proc. IWK Ilmenau, pp. 382-393, 2010.

[14] H.-M. Gross, C. Schröter, S. Müller, M. Volkhardt, E. Einhorn, A. Bley, Ch. Martin, T. Langer and M. Merten. I'll keep an Eye on You: Home Robot Companion for Elderly People with Cognitive Impairment. In IEEE Int. Conf. on SMC, pp. 2481-2488, 2011.

[15] D. Fox, W. Burgard and S. Thrun. The Dynamic Window Approach to Collision Avoidance. In IEEE Robotics & Automation Magazine,'97.

[16] J. Kessler, A. Scheidig, H.-M. Gross. Approaching a Person in a Socially Acceptable Manner Using Expanding Random Trees. In Proc. 5th European Conference on Mobile Robots, pp. 95-100, 2011

[17] M. Volkhardt, S. Mueller, Ch. Schroeter, and H.-M. Gross. Playing Hide and Seek with a Mobile Companion Robot. In Proc. 11th IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids), pp. 40–46, 2011.

[18] K.O. Arras, O.M. Mozos and W. Burgard. Using boosted features for the detection of people in 2d range data. In IEEE International Conference on Robotics and Automation, pp. 3402-3407, 2007.

[19] P. Viola and M. Jones. Robust real-time object detection. In International Journal of Computer Vision, vol. 57, no. 2, pp. 137-154, 2002.

[20] C. Martin, E. Schaffernicht, A. Scheidig and H.-M. Gross. Sensor Fusion using a Probabilistic Aggregation Scheme for People Detection and People Tracking. In Robotics and Autonomous Systems, vol. 54, issue 9, pp. 721-728, 2006.

[21] N. Dalal and B. Triggs. Histograms of Oriented Gradients for Human Detection. In Proc. IEEE Conf. on CVPR, pp. 886–893, 2005.

[22] V. Ferrari, M. Marin-Jimenez and A. Zisserman. Progressive search space reduction for human pose estimation. In IEEE Conference on Computer Vision and Pattern Recognition, pp. 1-8, 2008.