In Proc. of the IEEE Workshop on Advanced Robotics and its Social Impacts, Munich, Germany, May 2012

Benchmarking Mobile Manipulation in Everyday Environments

Jörg Stückler and Sven Behnke

Autonomous Intelligent Systems, University of Bonn, 53113 Bonn, Germany

Abstract— **In this paper, we discuss the RoboCup@Home league as a benchmark for service robot systems in everyday environments. The competition requires skills in mobile manipulation and human-robot interaction. Specifically, we detail the contributions of our team NimbRo, which won the RoboCup@Home competition in 2011. We demonstrated novel capabilities in the league such as real-time table-top segmentation, flexible grasp planning, real-time tracking of objects, and human-robot cooperative manipulation. We report on the experiences made with our robots at the competition.**

Fig. 1. Cognitive service robot *Cosero* in a supermarket and opening a bottle.

I. INTRODUCTION

Robot competitions such as the DARPA Grand and Urban Challenges [1], the European Land-Robot Trial (ELROB) [2], and—not the least—RoboCup [3], [4] provide a standardized test bed for robotic systems. They require participating teams to operate their robots in a common environment, outside their own lab, at a scheduled time. This makes it possible to directly compare the different approaches for robot construction, environment perception, and control.

The annual RoboCup competitions are well known for their soccer leagues. As a step towards applications, further leagues have been included such as the RoboCup Rescue league for robots supporting first responders and RoboCup@Home addressing service robot applications in everyday environments.

Since 2009, we participate in the RoboCup@Home league with great success. The competition fosters research on mobile manipulation and human-robot interaction. Since the application domain requires the integration of many capabilities, the approaches are integrated systems and benchmarking individual components becomes less suitable. Instead, benchmarking them can be conducted by demonstrating (and comparing) the performance and reliability of complete systems in a realistic setup and in an integrated way.

Integrated demonstrations of service robots in everyday environments are also performed by research groups in their own lab or at trade fairs, e.g. at TU Munich [5], at KIT [6], at DLR [7], at Willow Garage [8], and at Yaskawa [9]. While the technical achievements of these demonstrations are impressive, due to the isolated performances, it is hard to compare them directly to the work of other groups.

II. THE ROBOCUP@HOME LEAGUE

The RoboCup@Home league [10], [11] has been established in 2006 to foster the development and benchmarking of dexterous and versatile service robots that can operate safely in everyday scenarios. The robots have to show a wide variety of skills including object recognition and grasping, safe indoor navigation, and human-robot interaction (HRI). In 2011, 19 international teams competed in the @Home league. It is currently one of the strongest growing leagues in RoboCup.

The competition is organized into two preliminary rounds or *stages* and the Final. The stages consist of predefined test procedures as well as open demonstrations in which the teams can show what their robot can do best.

The lean rules in the RoboCup@Home league facilitate a variety of approaches. Some teams construct new and innovative robot hardware, while others resort to off-the-shelf hardware in order to focus on algorithmic problems. In the following, we discuss some approaches of teams that participated in the Robocup@Home competition in 2011.

The team WrightEagle [12] from the University of Science and Technology of China competes in the @Home league since 2009. They proposed a cognitive software architecture for their KeJia robot, which combines methods for natural language processing, reasoning, and task execution. In 2011, WrightEagle introduced the KeJia-2 robot platform that supports omni-directional driving and is equipped with two 7-DOF manipulators for human-like reach. In the competition, KeJia made popcorn in a microwave oven. For this demonstration, the robot had to press buttons and to open and close the microwave door.

The German team b-it-bots [13] from the Bonn-Rhein-Sieg University of Applied Sciences introduced their robot Jenny in this year's competition. Jenny consists of a modified Care-O-Bot 3 platform from Fraunhofer IPA with a 7-

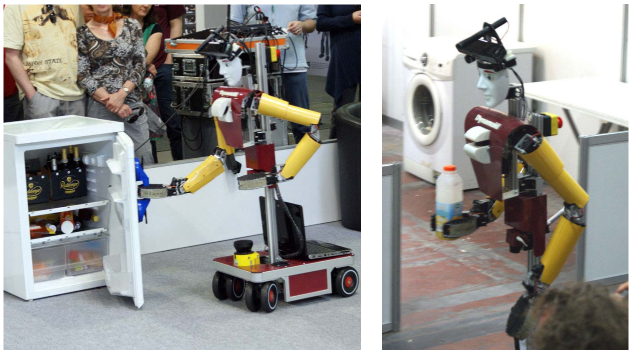

Fig. 2. Our service robot *Dynamaid* opens the fridge to fetch a drink.

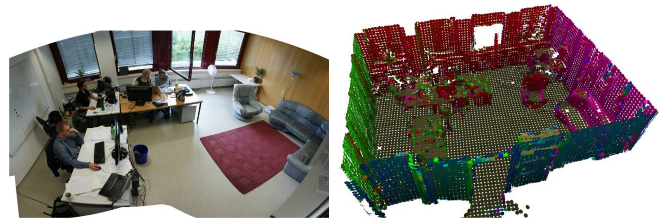

Fig. 3. Left: panorama image of an office. Right: 3D surfel map learned with our approach (surfel orientation coded by color).

DOF Kuka lightweight robot arm and a 3-finger hand from Schunk. They proposed means for object categorization, facial expression recognition, and interpretation of pointing gestures. Their control framework incorporates a deliberative layer implemented in a hierarchical task network.

The Australian team RobotAssist [14] from the ARC Centre of Excellence in Autonomous Systems competes with a mobile manipulation platform that is built from a Segway RMP 100 base and an Exact Dynamics iArm manipulator. For manipulator control, they apply an optimization method that finds collision-free arm configurations for the object to manipulate. RobotAssist also demonstrated person detection, identification, and social skills with their robot.

In this year's competition, our team NimbRo participated with Dynamaid and its successor, Cosero. In the tests, the robots showed their human-robot interaction and mobile manipulation capabilities. We introduced many new developments, like grasp planning to extend the range of graspable objects, real-time scene segmentation and object tracking, and human-robot cooperative manipulation of a table. Our performance was well received and has been awarded the first place in the RoboCup 2011 competition. In the following, we will detail our main contributions to the @Home league.

III. Robot Design

We focused the design of our robots Dynamaid and Cosero (s. Fig. 1 and Fig. 2) on typical requirements for autonomous operation in everyday tasks. While Cosero [15] still retains the light-weight design principles of Dynamaid [16], we improved its construction and appearance significantly and made it more precise and stronger actuated. Cosero's mobile base has a small footprint of 59×44 cm and drives omni-directionally. This allows the robot to maneuver through the narrow passages found in household environments. Its two anthropomorphic arms resemble average human body proportions and reaching capabilities. A yaw joint in the torso enlarges the workspace of the arms. In order to compensate for the missing torso pitch joint and legs, a linear actuator in the trunk can move the upper body vertically. This enables the robot to manipulate on similar heights like humans, even on the floor.

We constructed our robots from light-weight aluminum parts. All joints in the robots are driven by Robotis Dynamixel actuators. These design choices allow for a light-weight and inexpensive construction, compared to other domestic service robots. While each arm has a maximum payload of 1.5 kg and the drive has a maximum speed of 0.6 m/sec, the low weight (in total ca. 32 kg) requires only moderate actuator power.

Cosero perceives its environment with a variety of complementary sensors. The robot senses the environment in 3D with a Microsoft Kinect RGB-D camera in its pan-tilt head. For obstacle avoidance and tracking in farther ranges and larger field-of-views than the Kinect, the robot is equipped with multiple laser-range scanners. The sensor head also contains a shotgun microphone for speech recognition. By placing the microphone on the head, the robot points the microphone towards human users and at the same time directs its visual attention to them.

IV. Everyday Manipulation Skills of Our Robots

A. Mobile Manipulation

One significant part of the competition in the @Home league tests the mobile manipulation capabilities of the robots. They shall be able to fetch objects from various locations in the environment. To this end, the robot must navigate through the environment, perceive objects, and grasp them.

We implement 2D navigation with state-of-the-art methods. In static environments, Cosero localizes and plans paths in a 2D occupancy grid map. The main sensor for localization is the SICK S300 laser scanner on its mobile base. For 3D collision avoidance, we integrate measurements from any 3D sensing device, such as the tilting laser in the robot's chest.

We address several shortcomings of using 2D maps for localization and path planning by building 3D maps of the environment [17]. One problem of such 2D maps occurs in path planning, if untraversable obstacles cannot be perceived on the laser scanner's height. Localization with 2D lasers imposes further restrictions if dynamic objects occur, or the environment changes in the scan plane of the laser. Then, localization may fail since large parts of the measurements are not explained by the map.

We choose to represent the map in a 3D surfel grid which the robot acquires from multiple 3D scans of the environment. Fig. 3 demonstrates an example map generated with our approach. From the 3D maps, we extract 2D navigation maps by exploring the traversability of surfels.

We check for untraversable bumps between surfels and for obstacles within the robot's height range.

For localization, we developed an efficient Monte Carlo method that can incorporate full 3D scans as well as 2D scans. When used with 3D scans, we extract surfels from the scans and evaluate the observation likelihood. From 2D scans, we extract line segments and associate them with surfels in the map. Localization in 3D maps is specifically useful in crowded environments. The robot can then leverage measurements above the height of people to localize at the static parts of the environment. More general, by representing planar surface elements in the map, we can also rely for localization mainly on planar structures, as they more likely occur in static environment parts. For further details please refer to [17].

Typically, in mobile manipulation the robot estimates its pose in reference to the walls, objects, and persons. For example, when the robot grasps an object from a table, it first approaches the table roughly within the reference frame of the static map. Then, it adjusts in height and distance to the table. Finally, it aligns itself to bring the object into the workspace of its arms.

Cosero grasps objects on horizontal surfaces like the floor, tables and shelves in a height range from ca. 0 m to 1 m. It carries the object and hands it to human users. We also developed solutions to pour-out containers, to place objects on horizontal surfaces, to dispose objects in containers, to grasp objects from the floor, and to receive objects from users.

B. Compliance Control

From differential inverse kinematics, we derived a method to limit the torque of the joints depending on how much they contribute to the achievement of the motion in task-space [18]. Our approach not only allows to adjust compliance in the null-space of the motion, but also in the individual dimensions in task-space. This is very useful when only specific dimensions in task-space shall be controlled in a compliant way.

We applied compliant control to the opening and closing of doors that can be moved without the handling of an unlocking mechanism. Refrigerators or cabinets are commonly equipped with magnetically locked doors that can be pulled open without special manipulation of the handle. See Fig. 2 for an example. Several approaches exist to manipulate doors when no precise articulation model is known ([19], [20]). Our approach does not require feedback from force or tactile sensors.

To open a door, our robot drives in front of it, detects the door handle with the torso laser, approaches the handle, and grasps it. The drive moves backward while the gripper moves to a position to the side of the robot in which the opening angle of the door is sufficiently large to approach the open fridge or cabinet. The gripper follows the motion of the door handle through compliance in the lateral and the yaw directions. The robot moves backward until the gripper reaches its target position. For closing a door, the robot has to approach the open door leaf, grasp the handle, and move forward while it holds the handle at its initial grasping pose relative to the robot. When the arm is pulled away from this pose by the constraining motion of the door leaf, the drive corrects for the motion to keep the handle at its initial pose relative to the robot. The closing of the door can be detected when the arm is pushed back towards the robot.

(a) (b) (c)

Fig. 4. Table-top segmentation. (a) Example setting. (b) Raw point cloud from Kinect with RGB information. (c) Each detected object is marked with a distinct color.

C. Real-Time Table-Top Segmentation

In household environments, objects are frequently located on planar surfaces such as tables. We therefore base our object detection pipeline on fast planar segmentation of the depth images of the Kinect [15]. Fig. 4 shows an exemplary result of our approach in a table-top scene. Our approach processes depth images with a resolution of 160×120 at frame rates of approx. 16 Hz on the robot's main computer. This enables our system to extract information about the objects in a scene with a very low latency for further decision-making and planning stages. For object identification, we utilize texture and color information [16].

In order to process the depth images efficiently, we combine rapid normal estimation [21] with fast segmentation techniques. The normal estimation method utilizes integral images to estimate surface normals in a fixed image neighborhood in constant time. Overall, the runtime complexity is linear in the number of pixels for which normals are calculated. Since we search for horizontal support planes, we find all points with vertical normals. We segment these points into planes using RANSAC [22]. We find the objects by clustering the points above the convex hull of the support plane points.

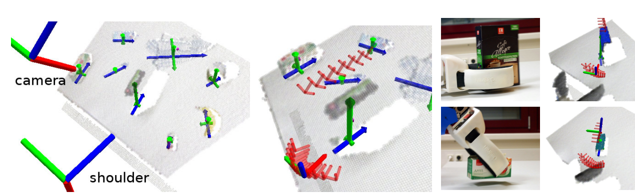

(a) (b) (c)

Fig. 5. Grasp planning. (a) Object shape properties. The arrows mark the principal axes of the object. (b) We rank feasible, collision-free grasps (red, size prop. to score) and select the most appropriate one (large, RGB-coded). (c) Example grasps on box-shaped objects.

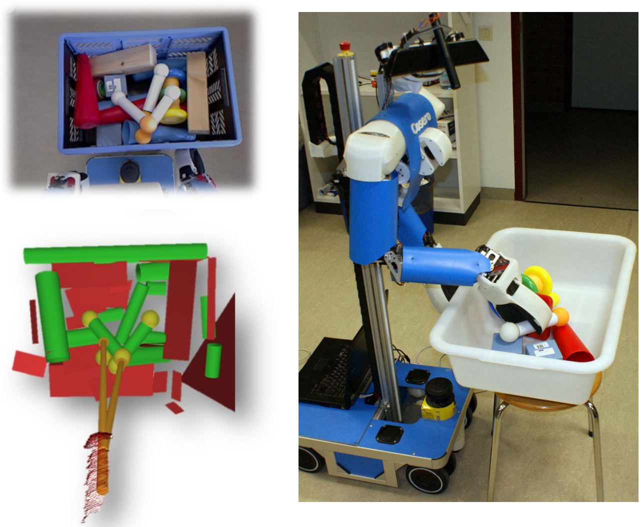

Fig. 6. Left: Objects composed from shape primitives recognized in a transport box. Right: Cosero executes a planned collision-free grasp on an object in a box.

D. Efficient Grasp Planning

We investigate grasp planning to enable our robots to grasp objects that they typically encounter in RoboCup. In order to grasp objects flexibly from shelves and in complex scenes, we consider obstructions by obstacles [15]. In our approach, we assume that the object is rigid and symmetric along the planes spanned by the principal axes of the object, e. g., cylindrical or box-shaped objects. We found that our approach also frequently yields stable grasps when an object violates these assumptions. Fig. 5 illustrates the main steps in our grasp planning pipeline and shows example grasps.

We consider two kinds of grasps: A side-grasp that approaches the object horizontally and grasps the object along the vertical axis in a power grip. The complementary top-grasp approaches the object from the top and grasps it with the finger tips along horizontal orientations. Our approach extracts the object's principle axes in the horizontal plane and its height. We sample pre-grasp postures for top- and side-grasps which we examine for feasibility under kinematic and collision constraints.

In on-going work, we advance our method to operate in more general scenes and to be less restrictive in collision checking (s. Fig.6). We apply the method by Schnabel et al. [23], [24] and Berner et al. [25] to detect objects composed from shape primitives such as cylinders, spheres, and cones in point clouds. On the shape primitives, we sample pre-grasp postures and identify feasible collision-free grasps. For fast grasp planning, we successively prune infeasible grasps in stages of ascending run-time complexity. Finally, we efficiently plan collision-free reaching motions using a multi-resolution sampling-based planner [26]. Our method is suitable for recognizing and grasping objects in complex scenes such as object piles.

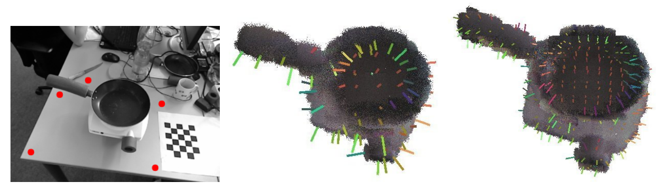

(a) (b) (c)

Fig. 7. Learning object models. (a) During training the user selects points (red dots) to form a convex hull around the object. (b) Color and shape distribution modeled at 2.5 cm resolution. Lines indicate surface normals (color-coded by orientation). (c) Color and shape distribution modeled at 5 cm resolution. Lines indicate surface normals (color-coded by orientation).

E. Real-Time Object Tracking

When a robot interacts with objects, it has to estimate its pose with respect to the objects. Frequently, the localization of the object in a map is not precise enough for this purpose. For example, the place of many household objects such as tables or chairs is subject to frequent changes. The robot must hence be able to detect the object in its current sensor view and estimate the relative pose of the object.

We developed methods for real-time tracking of objects with RGB-D cameras [27]. We train full-view multi-resolution surfel maps of objects (s. Fig. 7) and track these maps in RGB-D images in real-time.

Our maps represent the normal distribution of points including their color in voxels at multiple resolutions using octrees.

We register these maps to the object map with an efficient multi-resolution strategy. Instead of comparing the image pixel-wise to the map, we build multi-resolution surfel maps m_s with color information from new RGB-D images and register these maps to the object map m_m with an efficient multi-resolution strategy. We associate each node in the image map to its corresponding node in the object map using fast nearest-neighbor look-up. We measure the matching likelihood

$$p(m_s|x, m_m) = \prod_{(i,j)\in\mathcal{C}} p(s_{s,i}|x, s_{m,j}) \tag{1}$$

for the surfel correspondences, where $\mathcal{C}$ is the set of surfel correspondences between the maps, and $s_{s,i} = (\mu_{s,i}, \Sigma_{s,i})$ and $s_{m,j} = (\mu_{m,j}, \Sigma_{m,j})$ are corresponding surfels.. We iteratively optimize this likelihood to find the most likely transformation between the maps. In order to cope with illumination changes, we ignore minor luminance and color differences.

The observation likelihood of a surfel match is the difference of the surfels under their normal distributions,

$$\begin{aligned} p(s_s|x, s_m) &= \mathcal{N}\left(d(s_m, s_s, x); 0, \Sigma(s_m, s_s, x)\right), \\ &= \mathcal{N}\left(\mu_m - T(x)\mu_s; 0, \Sigma_m + R(x)\Sigma_s R(x)^T\right), \end{aligned} \tag{2}$$

where $T(x)$ is a transformation matrix that rotates and translates the spatial dimensions according to the pose x and $R(x)$ is the corresponding rotation matrix.

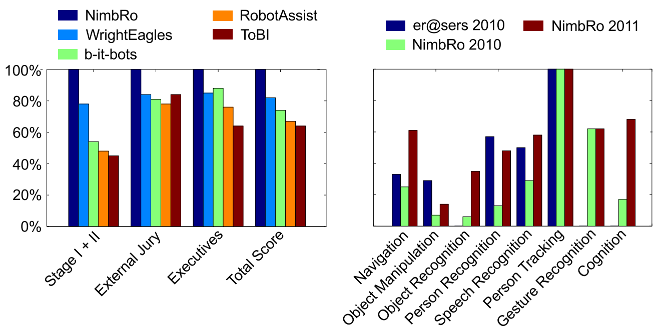

Fig. 8. Left: Relative scores per finalists 2011. Right: Reached scores in the predefined test procedures per functionality.

We associate surfels between maps using efficient nearest neighbor look-ups in the octree. In order to determine the correspondences between surfels in both maps, we apply a coarse-to-fine strategy that selects the finest resolution possible. We only establish a correspondence, if the surfels also match in the color cues. Our association strategy not only saves redundant comparisons on coarse resolution. It also allows to match surface elements at coarser scales if shape and color cannot be matched on finer resolutions. By this, our method allows the object to be tracked from a wide range of distances.

F. Human-Robot Cooperative Manipulation

We study physical interaction between a human user and a robot in a cooperative manipulation task [27]. In our scenario, the human and the robot cooperatively carry a large object, i. e., a table. For the successful performance of this task, the robot must keep track of the object and the actions of the human.

In order to accurately approach the table, the robot tracks the 6 DoF pose of the table in real-time. The user can then lift and lower the table, which the robot simply perceives through the motion of the registered table model. Once the table is lifted, the robot lifts the table as well and sets its arms compliant in the horizontal plane and in vertical orientation. This enables the human to move the robot arms through the table. The robot follows this motion until the human puts the table down again.

V. Performance at RoboCup 2011

With Dynamaid and Cosero, we competed in the RoboCup @Home 2011 competition in Istanbul. Our robots participated in all tests of Stages I and II, and performed very well. We accumulated the highest score of all 19 teams in both stages. Our final demonstration was also awarded the best score such that we achieved the first place in the competition.

A. Stage I

Stage I begins with the *Robot Inspection and Poster Session* test. Our robots registered themselves at the registration desk, while we presented our work to the leaders of other teams in a poster session. Overall, we received the highest score in this test. In the *Follow Me* test, Cosero met a previously unknown person and followed him reliably through an unknown environment. Cosero could show, that it can distinguish this person from others, and that it recognizes stop gestures. In *Who Is Who*, two previously unknown persons introduced themselves to Cosero. Later in the test, our robot found one of the previously unknown persons, two team members, and one unknown person and recognized their identity correctly. The *Open Challenge* allows the teams to show their research in self-defined tasks. Cosero fetched a bottle of milk, opened it, and poured it into a cereal bowl. Then, Cosero grasped a spoon using our approach to grasp planning and placed it next to the bowl. Cosero understood a complex command partially and went to a correct place in the *General Purpose Service Robot I* test. In the *Go Get It!* test, Cosero found a correct object and delivered it. After Stage I, we were leading the competition.

B. Stage II

In Stage II, Cosero participated in the *Shopping Mall* test. It learned a map of a previously unknown shopping mall and navigated to a shown location. Taking a shopping order was hindered by speech-recognition failures in the unknown acoustic environment. In the *General Purpose Service Robot II* test, Cosero first understood a partially specified command and asked questions to obtain missing information about the object to grasp and the location of the object. It executed the task successfully. In the second part of the test, it worked on a task with erroneous information. It detected that the ordered object was not at the specified location, went back to the user, and reported the error. In the *Demo Challenge*, we demonstrated pointing gestures by showing the robot in which baskets to put colored and white laundry. The robot cleaned the apartment, picked white laundry from the floor, and put it into the correct basket. Afterwards it cleaned up the appartment and picked objects from a table. The technical committee awarded us the highest score.

Fig. 8 summarizes the scores achieved for individual functionalities as proposed in [28]. Note that due to the sequential nature of the predefined test procedures, in some tests our robots did not reach specific sub-tasks. For instance, in *Enhanced Who Is Who* or *Shopping Mall*, our system had difficulties to understand the orders by the human user and, hence, did not have the chance to gain score for object manipulation. The results demonstrate that we were able to improve most functionalities compared to 2010 and achieved well in developing a balanced domestic service robot system. Overall, we reached the Final with 8,462 points, followed by Wright Eagle from China with 6,625 points.

C. Final

In the Final, we demonstrated the cooperative carrying of a table by Cosero and a human user (s. Fig. 9). Then, a user showed Cosero where it finds a bottle of dough to make an omelet. Our robot went to the cooking plate to switch it on. It succeeded partially in turning the plate on. Then, it drove to the location of the dough and grasped it. At the cooking plate, it opened the bottle and poured it into the pan. We applied

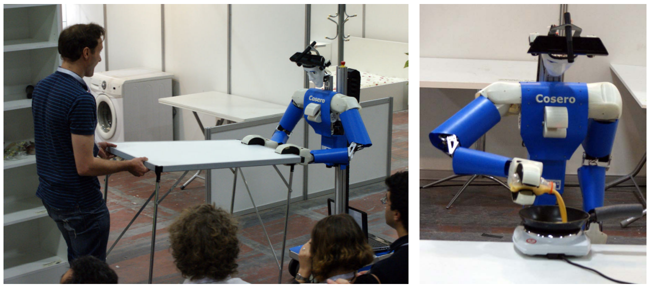

Fig. 9. *Cosero* cooperatively carries a table with a user and bakes omelet during the RoboCup@Home Final 2011 in Istanbul.

our real-time object tracking method in order to approach the cooking plate. Meanwhile, Dynamaid opened a refrigerator and grasped a bottle of orange juice out of it. It placed the orange juice on the breakfast table. Fig. 8 shows the relative scores of the finalists in 2011. Our performance received the best score by the juries that consisted of members of the executive committee and external judges from science and the media.

VI. Conclusion

The RoboCup@Home league is a competition for service robots in everyday environments. It benchmarks mobile manipulation and HRI capabilities of integrated robotic systems.

As a specific example of approaches in the RoboCup @Home league, we presented the contributions of our winning team NimbRo. In this paper, we detailed our methods for real-time scene segmentation, object tracking, and human-robot cooperative manipulation. In the tests in Stages I and II, our robots Cosero and Dynamaid performed very well. Our advanced mobile manipulation and HRI skills have been well received by juries in the open demonstrations and the Finale.

In future work, we aim to further advance the versatility of the mobile manipulation and human-robot interaction skills of our robots. The learning of models of arbitrary objects and the real-time tracking of these models is one step in this direction. Equally important, we are working to improve the perception of persons and the interpretation of their actions. We also plan to remove the necessities to adapt the tools that the robot uses to its current end-effectors. In order to improve the manipulation skills of our robots, we will improve the design of the grippers. We plan to construct thinner fingers with touch sensors. Then, we can devise new methods to grasp smaller objects or to use smaller tools.

References

[1] DARPA Grand Challange. http://archive.darpa.mil/grandchallenge.

[2] ELROB: The European Robot Trial. http://www.elrob.org.

[3] The RoboCup Federation. http://www.robocup.org.

[4] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, and E. Osawa. RoboCup: The Robot World Cup Initiative. In *Proc. of 1st Int. Conf. on Autonomous Agents (AGENTS)*, 1997.

[5] M. Beetz, U. Klank, I. Kresse, A. Maldonado, L. Msenlechner, D. Pangercic, T. Rühr, and M. Tenorth. Robotic roommates making pancakes. In *11th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, Bled, Slovenia, 2011.

[6] Z. Xue, S.W. Ruehl, A. Hermann, T. Kerscher, and R. Dillmann. An autonomous ice-cream serving robot. In *Proc. of Int. Conference on Robotics and Automation (ICRA)*, pages 3451–3452. IEEE, 2011.

[7] B. Bäuml, F. Schmidt, T. Wimböck, O. Birbach, A. Dietrich, M. Fuchs, W. Friedl, U. Frese, C. Borst, M. Grebenstein, O. Eiberger, and G. Hirzinger. Catching flying balls and preparing coffee: Humanoid Rollin'Justin performs dynamic and sensitive tasks. In *Proc. of Int. Conf. on Robotics and Automation (ICRA)*, pages 3443–3444, 2011.

[8] J. Bohren, R.B. Rusu, E.G. Jones, E. Marder-Eppstein, C. Pantofaru, M. Wise, L. Mösenlechner, W. Meeussen, and S. Holzer. Towards autonomous robotic butlers: Lessons learned with the PR2. In *Proc. of ICRA*, pages 5568–5575, 2011.

[9] M. Adachi, Z. Zhang, P. Chen, and Q. Cao. An RTC-based intelligent programming system for service robots. In *Intelligent Robotics and Applications*, volume 6424 of *LNCS*, pages 226–236. Springer, 2010.

[10] T. van der Zant and T. Wisspeintner. RoboCup X: A proposal for a new league where robocup goes real world. In *RoboCup 2005: Robot Soccer World Cup IX*, LNCS 4020, pages 166–172, 2005.

[11] T. Wisspeintner, T. van der Zant, L. Iocchi, and S. Schiffer. RoboCup@Home: Scientific Competition and Benchmarking for Domestic Service Robots. *Interaction Studies*, 10(3):393–428, 2009.

[12] X. Chen, G. Jin, J. Ji, F. Wang, and Jiongkun Xie. KeJia project: Towards integrated intelligence for service robots. In *RoboCup@Home League Team Descriptions, Istanbul*, 2011.

[13] F. Hegger, C. Müller, Z. Jin, J. A. lvarez Ruiz, G. Giorgana, N. Hochgeschwender, M. Reckhaus, J. Paulus, P. Ploeger, and G. K. Kraetzschmar. The b-it-bots RoboCup@Home 2011 team description paper. In *RoboCup@Home League Team Descriptions, Istanbul*, 2011.

[14] A. Alempijevic, S. Carnian, D. Egan-Wyer, G. Dissanayake, R. Fitch, B. Hengst, D. Hordern, N. Kirchner, M. Koob, M. Pagnucco, C. Sammut, and A. Virgona. RobotAssist - RoboCup@Home 2011 team description paper. In *RoboCup@Home League Team Descriptions, Istanbul*, 2011.

[15] J. Stückler, R. Steffens, D. Holz, and S. Behnke. Real-time 3D perception and efficient grasp planning for everyday manipulation tasks. In *Proc. of the Europ. Conf. on Mobile Robots (EMCR)*, 2011.

[16] J. Stückler and S. Behnke. Integrating indoor mobility, object manipulation, and intuitive interaction for domestic service tasks. In *Proc. of the IEEE Int. Conf. on Humanoid Robots (Humanoids)*, 2009.

[17] J. Kläß, J. Stückler, and S. Behnke. Efficient mobile robot navigation using 3D surfel grid maps. In *Proc. of 7th German Conf. on Robotics (ROBOTIK)*, 2012.

[18] J. Stückler and S. Behnke. Compliant task-space control with back-drivable servo actuators. In *Proc. of the 15th RoboCup International Symposium*, 2011.

[19] G. Niemeyer and J.-J. E. Slotine. A simple strategy for opening an unknown door. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1997.

[20] A. Jain and C.C. Kemp. Pulling open doors and drawers: Coordinating an omni-directional base and a compliant arm with equilibrium point control. In *Proc. Int. Conf. on Robotics and Automation (ICRA)*, 2010.

[21] D. Holz, S. Holzer, R. B. Rusu, and S. Behnke. Real-time plane segmentation using RGB-D cameras. In *Proc. of the 15th RoboCup Int. Symp.*, Turkey, 2011.

[22] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Comm. of the ACM*, 24(6):381–395, 1981.

[23] R. Schnabel, R. Wahl, and R. Klein. Efficient RANSAC for point-cloud shape detection. *Computer Graphics Forum*, 26(2):214–226, 2007.

[24] R. Schnabel, R. Wessel, R. Wahl, and R. Klein. Shape recognition in 3D point-clouds. In *The 16th Int. Conference in Central Europe on Computer Graphics, Visualization, and Computer Vision*, 2008.

[25] A. Berner, M. Bokeloh, M. Wand, A. Schilling, and H.-P. Seidel. A graph-based approach to symmetry detection. In *Symposium on Volume and Point-Based Graphics*, pages 1–8, 2008.

[26] I.A. Sucan and L.E. Kavraki. Kinodynamic motion planning by interior-exterior cell exploration. In *Proc. of 8th WS on Algorithmic Foundations of Robotics*, volume 57, pages 449–464, 2009.

[27] J. Stückler and S. Behnke. Following human guidance to cooperatively carry a large object. In *Proc. of the 11th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2011.

[28] T. Wisspeintner, T. van der Zant, L. Iocchi, and S. Schiffer. RoboCup@Home: Scientific competition and benchmarking for domestic service robots. *Interaction Studies*, 10(3):392–426, 2009.