IEEE Robotics & Automation Magazine 19(2):34-42, June 2012.

Demonstrating Everyday Manipulation Skills in RoboCup@Home

Jörg Stückler, Dirk Holz *Member, IEEE*, and Sven Behnke *Member, IEEE*

***Abstract*—The RoboCup@Home league is a benchmark for domestic service robot systems. It evaluates approaches to mobile manipulation and human-robot interaction by testing integrated systems. In this article, we detail the contributions of our team NimbRo, with which we won the RoboCup@Home competition in 2011. We demonstrated novel capabilities in the league such as real-time tabletop segmentation, flexible grasp planning, and real-time tracking of objects. We also describe our approach to human-robot cooperative manipulation using compliant control. We report on the use of our approaches and the performance of our robots at RoboCup 2011.**

***Index Terms*—RoboCup@Home, mobile manipulation, real-time scene perception, benchmarks in robotics.**

I. Introduction

As benchmarking robotics research is inherently difficult, robot competitions are increasingly popular. They bring together researchers, students, and enthusiasts in the pursuit of a technological challenge. Prominent examples for such competitions include the DARPA Grand and Urban Challenges [1], the International Aerial Robotics Competition (IARC) [2], the European Land-Robot Trial (ELROB) [3], and—not the least—RoboCup [4], [5].

Such competitions provide a standardized test bed for different robotic systems. All participating teams are forced to operate their robots outside their own lab in an uncontrolled environment at a scheduled time. This makes it possible to directly compare the different approaches for robot construction, environment perception, and control.

While the annual RoboCup competitions are best known for their soccer leagues, they also feature two leagues in other domains—the RoboCup Rescue league for robots supporting first responders and RoboCup@Home addressing service robot applications in domestic environments.

In RoboCup@Home, different disciplines of robotics research such as, for instance, mobile manipulation and human-robot interaction are tightly coupled. That is, approaches are integrated systems and benchmarking individual components becomes less suitable. Instead, benchmarking them is conducted by demonstrating (and comparing) the performance and reliability of complete systems in a realistic setup and in an integrated way.

In this article, we present the contributions of our team NimbRo to the RoboCup@Home league. We describe the challenges in this league and detail our approaches to the competition. Our team achieved first place in the 2011 competition.

All authors are with the Autonomous Intelligent Systems Group, Computer Science Institute VI, University of Bonn, 53113 Bonn, Germany, e-mail: stueckler@ais.uni-bonn.de.

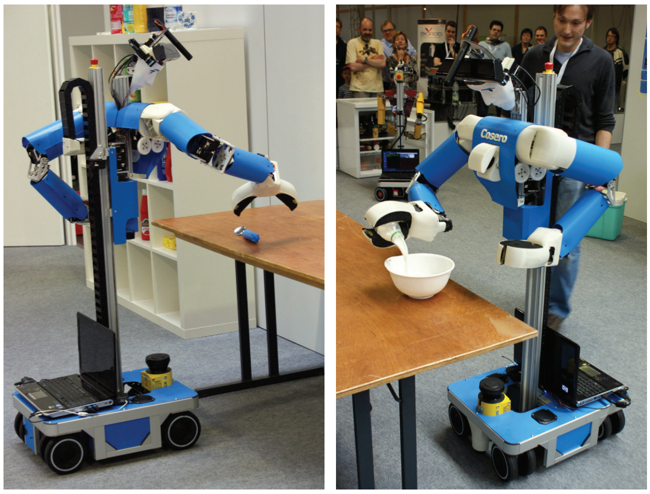

Fig. 1. Cognitive service robot *Cosero* grasps a spoon and pours milk into a bowl of cereals at RoboCup GermanOpen 2011.

While we successfully participated in many standard tests, we also demonstrated novel capabilities in the league such as real-time tabletop segmentation, flexible grasp planning, and real-time tracking of objects. We also describe our approach to human-robot cooperative manipulation using compliant control.

II. The RoboCup@Home League

The RoboCup@Home league [6], [7] was 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.

A. Competition Design

The competition is organized into two preliminary rounds or *stages* and a final [8]. The stages consist of predefined test procedures as well as open demonstrations.

The predefined tests include skills for domestic service robots that must be solved with state-of-the-art approaches. The time to accomplish the tests is limited, which forces the teams to implement time-efficient approaches. During the tests, the robots must operate autonomously. Helping with physical interaction or remote control is not allowed. The rules also include extra scores for specific skills that require solutions to research questions, rewarding scientific solutions that go beyond the fulfillment of the basic requirements of a test. In the open demonstrations, the teams can choose their own task for the robot in order to demonstrate results of their own research.

The top 50% of the teams (w.r.t. score) after the first stage advance to the second stage where they have to perform more complex tasks. The top 50% of the teams (w.r.t. score) after the second stage (including the points scored in the first stage) further advance to the Final which is conducted as an open demonstration.

While the rules and the tests are announced several months prior to the competition, the details of the competition environment are not known to the participants in advance. During the first two days of the competition, the teams can map the competition arena, which resembles an apartment, and train object recognition on a set of about 20 smaller objects which are used as known objects with names throughout the recognition and manipulation tests. The arena is subject to minor and major changes during the competition and also contains previously unknown objects.

B. Performance Evaluation

In the predefined tests, each sub-task is assigned a certain number of points which are awarded upon successful completion. This allows for an objective evaluation of the overall system's performance, as well as for the assessment of individual components.

The performances of the teams in the open demonstrations vary greatly and are hence harder to compare. Open demonstrations are evaluated by juries for their technical and scientific merits. In order to provide a fair assessment, these juries are formed from leaders of other teams or from members of the technical and executive committees of the league. The jury in the Final is formed from members of the league's executive committee and distinguished external representatives of science, industry, and the media.

The juries evaluate the teams by specific criteria that are defined in the rules of the competition. In the Final, for example, the external jury assesses originality and presentation, usability of human-robot interaction, difficulty and success, and relevance for daily life.

C. Tests and Skills

The tests in the RoboCup@Home league are designed to reflect (and test for) the large diversity of problems addressed in service robotics research.

1) Tests in Stage I: All teams participate in the first stage which tests basic mobile manipulation and HRI capabilities. In the *Robot Inspection and Poster Session* test, the robots have to navigate to a registration desk, introduce themselves, and get inspected by the league's technical committee. Meanwhile, the team gives a poster presentation that is evaluated by the leaders of the other teams. In *Follow Me*, the robots must demonstrate person tracking and recognition capabilities in an unknown environment. The robot is guided by a previously unknown user who can command the robot either by speech or by gestures. At several checkpoints, the robustness of the approaches is tested by applying different disturbances. Mobile manipulation and HRI capabilities have to be integrated for *GoGetIt*. Here, the robot has to retrieve the correct object among others from a room. The room is specified to the robot by a human user using speech input. Person detection and recognition in the home environment is tested within *WhoIsWho*. The robot has to learn the identity of two persons and must later find the persons among others in a different room. In the *General Purpose Service Robot I* test, the robots must understand and act according to complex speech commands that consist of three sub-tasks such as moving to a location, retrieving a specific object, and bringing it back to the user. The last test in this stage is the *Open Challenge* in which the teams can demonstrate their system in a five minute slot.

2) Tests in Stage II: The teams that advance to the second stage are tested in more complex scenarios. *Enhanced WhoIsWho* extends *WhoIsWho* towards a robotic butler scenario. A user tells the robot to bring beverages to three out of five persons. The robot has to fetch the beverages and to deliver them to the correct person. Again, the robot is introduced to the persons at the beginning. This test has to be solved within 10 min. In *General Purpose Service Robot II*, commands with missing or erroneous information are given to the robot. The robot has to ask for missing information, or, if it detects erroneous task specifications during the execution of the task, it must react accordingly, report back to the user, and propose alternative solutions. *Shopping Mall* tests the abilities of the robots to operate in previously unknown environments. A human user guides the robot through a real shopping mall and shows it the location of several objects. Afterwards, the robot must fetch a subset of the objects as specified by the human user. Stage II concludes with the *Demo Challenge*. This 7 min open demonstration follows a theme that is defined prior to the competition. In 2011, the theme was "Cleaning the House".

III. Related Research on Integrated Systems

The lean rules in the RoboCup@Home league facilitate diverse approaches. Some teams construct new and innovative robot hardware, while others resort to off-the-shelf hardware in order to focus on algorithmic problems.

The Chinese team WrightEagle [9] has competed in the @Home league since 2009. In 2011, they introduced the KeJia-2 robot platform that supports omni-directional driving and is equipped with two 7-DOF manipulators for human-like reach, similar to our robots. In the competition, KeJia made popcorn in a microwave oven. For this demonstration, the robot had to press buttons to open and close the microwave door.

The German team b-it-bots [10] introduced their robot Jenny in the 2011 competition. Jenny consists of a modified Care-O-Bot 3 platform from Fraunhofer IPA with a 7-DOF Kuka lightweight robot arm and a three-finger Schunk hand.

The Australian team RobotAssist [11] competes with a robot that combines a Segway RMP 100 base with an Exact

Fig. 2. Domestic service robot *Dynamaid* opens and closes the fridge during the RoboCup@Home Final 2010 in Singapore.

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.

Besides RoboCup@Home, many research groups currently develop integrated systems for mobile manipulation in everyday environments. Demonstrations of these systems are performed in isolated settings in labs or at trade fairs.

A prominent example is the Personal Robot 2 (PR2) developed by Willow Garage. Bohren et al. [12] demonstrate an application in which a PR2 fetches drinks from a refrigerator and delivers them to human users. Both the drink order and the location at which it has to be delivered are specified by the user in a web form. In Beetz et al. [13] a PR2 and a custom-built robot cooperatively prepare pancakes. In the healthcare domain, Jain and Kemp [14] present EL-E, a mobile manipulator that assists motor impaired patients by performing pick and place operations to retrieve objects. Srinivasa et al. [15] combine object search and retrieval in different demonstrations in their lab. Their autonomous service robot HERB navigates around a kitchen, searches for mugs and brings them back to the kitchen sink. Xue et al. [16] demonstrated grasping and handling of ice cream scoops with a two-armed robot standing at a fixed position. The DLR robot Rollin' Justin prepared coffee in a pad machine [17]. The robot grasped coffee pads and inserted them into the coffee machine, which involved opening and closing the pad drawer.

In the RoboCup 2011 competition, our team NimbRo participated with the robot 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 carrying of a table.

IV. System Overview

A. Robot Design

We focused the design of our robots Dynamaid [18] and Cosero [19] (see Figs. 1, 2) on typical requirements for autonomous operation in everyday environments. While Cosero still retains the light-weight design principles of Dynamaid, 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 omnidirectionally. 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 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 of Cosero 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. The robot's main computer is a quadcore notebook with an Intel i7-Q720 processor.

Both robots perceive their environment with a variety of complementary sensors. The robots sense the environment in 3D with a Microsoft Kinect RGB-D camera in their pan-tilt head. For obstacle avoidance and tracking in farther ranges and larger field-of-views than the Kinect, the robots are equipped with multiple laser-range scanners, from which one can be pitched and one can be rolled. The sensor head of the robots also contains a shotgun microphone for speech recognition. By placing the microphone on the head, the robots point the microphone towards human users and at the same time direct their visual attention to them.

B. Perception and Control Framework

The autonomous behavior of our robots is generated in a modular control architecture. We employ the inter process communication infrastructure and tools of the Robot Operating System (ROS) [20].

We implement task execution, mobile manipulation, and motion control in hierarchical finite state machines. The task execution level is interweaved with human-robot interaction modalities. For example, we support the parsing of natural language to understand and execute complex commands.

Tasks that involve mobile manipulation trigger and parametrize sub-processes on a second layer of finite state machines. These processes configure the perception of objects and persons, and they execute motions of body parts of the robot. The motions themselves are controlled on the lowest layer of the hierarchy and can also adapt to sensory measurements.

V. Everyday Manipulation Skills

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

We implement navigation with state-of-the-art methods. Cosero localizes and plans paths in a 2D occupancy grid map of the environment ([21], [22], [23]). For 3D collision avoidance, we integrate measurements from any 3D sensing device, such as the tilting laser in the robot's chest. Due to the limited on-board computing power of our robots, we focused on efficient and light-weight implementations.

In mobile manipulation, the robot typically estimates its pose in reference to the walls, objects, or persons. For example, when the robot grasps an object from a table, it first approaches the table roughly within the reference frame of a 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.

Our robots grasp objects on horizontal surfaces like the floor, tables, and shelves in a height range from the floor to ca. 1 m. They carry the objects and hand them to human users. We also developed solutions to pour-out containers, to place objects on horizontal surfaces, to dispose objects in containers, and to receive objects from users. We implemented these capabilities by parametrized motion primitives and also account for collisions during grasping motions.

A. 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 [24]. Our approach not only allows to adjust compliance in the null-space of the motion, but also in the individual dimensions of the 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 ([25], [26]). Our approach does not require feedback from force or tactile sensors. Instead, the actuators are back-drivable and measure the displacement due to external forces.

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. 3. Tabletop segmentation. (a) Example setting. (b) Raw colored point cloud from Kinect. (c) Each detected object is marked with a distinct color.

B. Real-Time Tabletop 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 [19]. Fig. 3 shows an exemplary result for a tabletop scene. Our approach processes depth images with a resolution of 160×120 at frame rates of approx. 20 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 [18].

Similar to Rusu et al. [27], we segment point clouds into objects on planar surfaces. In order to process the depth images efficiently, we combine rapid normal estimation with fast segmentation techniques. The normal estimation method utilizes integral images to estimate surface normals in a fixed image neighborhood in constant time [28]. 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 [29]. We find the objects by clustering the measurements above the convex hull of the points in the support plane.

C. Efficient Grasp Planning

We investigated 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 [19].

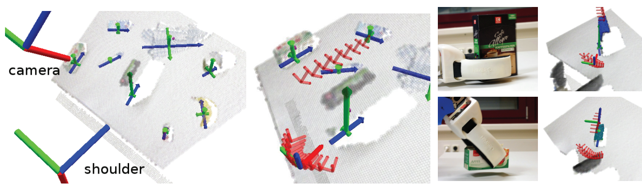

(a) (b) (c)

Fig. 4. 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.

Related approaches measure grasp quality, e. g., in the grasp wrench space [30], and virtually test grasps in physical simulation [31] in a time-costly process. We observed, however, that a well designed gripper, simple grasp strategies, and a compliant robot mechanism often suffice to grasp a large variety of household objects. Most related to our method is the approach by Hsiao et al. [32]. They use a time-consuming sampling-based motion planner to find collision-free reaching motions. In many situations, though, the direct reach towards the object is collision-free, or only few obstacles obstruct the motion. We thus apply parametrized motion primitives and take a conservative but efficient approach that checks simplified geometric constraints to detect collisions.

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. We found that our approach also frequently yields stable grasps when an object violates these assumptions. Fig. 4 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 detail, we consider the following feasibility criteria:

- *Grasp width.* We reject grasps, if the object's width orthogonal to the grasp direction does not fit into the gripper.
- *Object height.* Side-grasps are likely to fail if the object height is too small.
- *Reachability.* We do not consider grasps that are outside of the arm's workspace.
- *Collisions.* We check for collisions during the reaching and grasping motion.

The remaining grasps are ranked according to efficiency and robustness criteria:

- *Distance to object center.* We favor grasps with a smaller distance to the object center.
- *Grasp width.* We reward grasp widths closer to a preferred width (0.08 m).
- *Grasp orientation.* Preference is given to grasps with a smaller angle between the line towards the shoulder and the grasping direction.
- *Distance from robot.* We prefer grasps with a smaller distance to the shoulder.

The best grasp is selected and finally executed with a parametrized motion primitive.

D. Real-Time Object Tracking

The location of many household objects such as tables or chairs is subject to frequent changes. A robot must hence be able to detect objects in its current sensor view and estimate the relative pose of the objects.

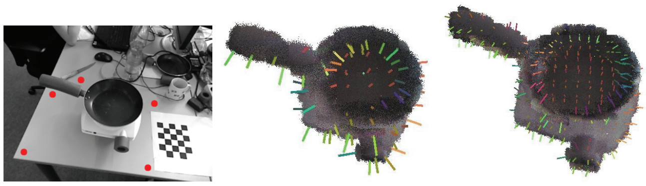

(a) (b) (c)

Fig. 5. 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 5 cm resolution. Lines indicate surface normals (color-coded by orientation). (c) Color and shape distribution modeled at 2.5 cm resolution.

We developed methods for real-time tracking of objects with RGB-D cameras [33]. We train full-view multi-resolution surfel maps of objects (see Fig. 5), and track these models in RGB-D images in real-time. Our method operates on 160×120 images at frame-rates of ca. 20 Hz on the robot's on-board computer.

Our maps represent the normal distribution of points including their color in voxels at multiple resolutions using octrees. Instead of comparing the image pixel-wise to the map, we build multi-resolution surfel maps with color information from new RGB-D images.

We register these maps to the object map with an efficient multi-resolution strategy. To this end, we measure the observation likelihood of the current image under the normal distributions of the surfels in both maps, and determine the most likely pose through optimization of this likelihood. In order to cope with illumination changes, we ignore minor luminance and color differences.

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 matches 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.

VI. Human-Robot Interaction

A service robot in everyday environments not only needs mobile manipulation abilities—it closely interacts with humans, even physically. This interaction should be natural and intuitive such that laymen can operate the robot and understand its actions.

In order to be aware of potential interaction partners, our robots detect and keep track of the persons in their surroundings [34]. Users can utter complex sentences to the robots, which the robots recognize and parse for semantics. Our robots also synthesize human-like speech. Furthermore, we equipped our robots with non-verbal communication cues. The robots can perform several gestures like pointing or waving. They can also perceive gestures such as pointing, showing of objects, or stop gestures [35] with the RGB-D camera.

Fig. 6. Left: *Cosero* and *Dynamaid* register themselves for the RoboCup@Home 2011 competition. Right: Cosero opens a bottle of milk during the *Open Challenge* at RoboCup 2011.

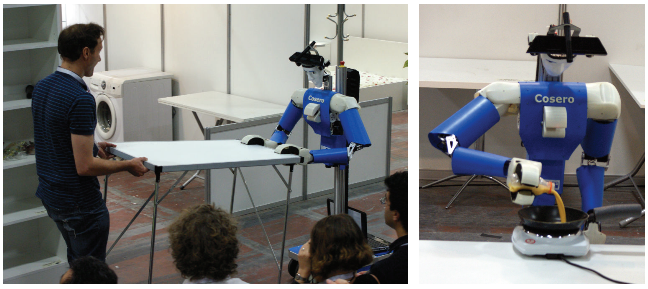

Fig. 7. *Cosero* cooperatively carries a table with a user and cooks omelet during the 2011 RoboCup@Home Final in Istanbul.

A. *Semantic Speech Interpretation*

We rely on the commercial Loquendo system for speech recognition and synthesis. Loquendo's speech recognition is grammar-based and speaker-independent. Its grammar definition allows rules to be tagged with semantic attributes. For instance, one can define keywords for actions or attributes like "unspecific" for location identifiers such as "room". When Loquendo recognizes a sentence that fits to the grammar, it provides the recognized set of rules together with a semantic parse tree. Our task execution module then interprets the resulting semantics and generates appropriate behavior.

B. *Human-Robot Cooperative Manipulation*

We compiled mobile manipulation, object perception, and human-robot interaction capabilities in a cooperative manipulation task [33]. In our scenario, the human and the robot cooperatively carry a table. For successful performance of this task, the robot must keep track of the table and the actions of the human. In order to accurately approach the table, the robot tracks its pose in real-time. The user can then lift and lower the table, which the robot simply perceives through the motion of the table. The robot follows the pulling and pushing on the table by the user through compliant control of its arms.

VII. Experiences at RoboCup 2011

With Dynamaid and Cosero, we competed in the RoboCup@Home 2011 competition in Istanbul. Our robots participated in all tests of Stage 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. Hence, we achieved the first place in the competition.

A. *Competition Performance*

In Stage I, Cosero and Dynamaid registered themselves in the *Robot Inspection and Poster Session* test, while we presented our work in a poster session. The robots generated speech and gestures and handed over the registration form. The leaders of the other teams awarded us the highest score in this test. In *Follow Me*, Cosero met a previously unknown person and followed him reliably through an unknown environment. Cosero could show, that it distinguishes this person from others, and that it recognizes stop gestures. In the *Who Is Who* test, two previously unknown persons introduced themselves to Cosero. Later in the test, our robot found one of the previously unknown persons, two members of our team, and one unknown person and recognized their identity correctly. In the *Open Challenge*, 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 speech command partially and went to the correct place in the *General Purpose Service Robot I* test. In *GoGetIt*, Cosero found a correct object and delivered it. After Stage I, we were leading the competition.

In the second stage, Cosero participated in *Shopping Mall*. It learned a map of a previously unknown area 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 an object and its location. 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 then cleaned the apartment, picked white laundry from the floor, and put it into the correct basket. It then picked carrots and tea boxes from a table. The objects could be chosen and placed by a jury member. The technical committee awarded us the highest score. We reached the Final with 8,462 points, followed by Wright Eagle from China with 6,625 points.

In the Final, we demonstrated the cooperative carrying of a table by Cosero and a human user (see Fig. 7). Then, a user showed Cosero where it finds a bottle of omelet mixture. 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 mixture and grasped it. At the cooking plate, it opened the bottle and poured it into the pan. We applied 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, which it then placed on the breakfast table. Our performance received the best score from the high-profile jury.

B. Lessons Learned

The experiences made at RoboCup 2011 clearly demonstrate our success in designing a balanced system that incorporates navigation, mobile manipulation, and intuitive human-robot interaction. The development of the system gave us many insights into the requirements and future steps towards complex domestic service scenarios.

Since the competition setting is unknown in advance, we have to develop methods that robustly work in a wide range of environments. We are also forced to implement means to adapt our approaches to new scenarios easily and in a fast way. For example, it is important to develop tools that allow maps, objects, and persons to be enrolled quickly. Such robust and fast-adaptable methods will be enablers for practical use.

In the typical manipulation scenarios that we encounter in the competition, our efficient grasping strategy seems more practical than traditional planning approaches w.r.t. time-efficiency and robustness in the presence of uncertainty. For complex manipulation settings such as grasping objects out of drawers and boxes, it will be necessary to develop efficient grasp and motion planning techniques that reason about uncertainties.

We have demonstrated that quite complex high-level behavior can be generated by semantic parsing of natural language and by a well designed hierarchical state-machine. It will be fruitful to push the complexity of the tasks with the versatility in skills. Then, new requirements will arise on reasoning capabilities for task execution and on semantic perception.

VIII. Conclusion

The RoboCup@Home league is a competition for service robots in domestic environments. It benchmarks mobile manipulation and HRI capabilities of integrated robotic systems. In this article, we presented the contributions of our winning team NimbRo. We detailed our methods for real-time scene segmentation, object tracking, and human-robot cooperative manipulation. In the pre-defined tests, we could demonstrate that our robots Cosero and Dynamaid solve mobile manipulation and HRI tasks with high reliability. Our advanced mobile manipulation and HRI skills have been well received by juries in the open demonstrations and the Final.

In future work, we aim to further advance the versatility of the skills of our robots. We constantly enhance our approaches to object and person perception. In order to extend 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 tools.

Acknowledgments

This research has been partially funded by the FP7 ICT-2007.2.2 project ECHORD (grant agreement 231143) experiment ActReMa. We thank the members of team NimbRo Kathrin Gräve, David Droeschel, Jochen Kläß, Michael Schreiber, and Ricarda Steffens for their dedicated efforts prior to and during the competition.

References

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

[2] "IARC: International Aerial Robotics Competition," http://iarc.angel-strike.com.

[3] "ELROB: The European Robot Trial," http://www.elrob.org.

[4] "The RoboCup Federation," http://www.robocup.org.

[5] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, and E. Osawa, "RoboCup: The robot world cup initiative," in *Proceedings of the 1st International Conference on Autonomous Agents*, New York, NY, USA, 1997, pp. 340–347.

[6] 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*, ser. LNCS 4020. Springer, 2006, pp. 166–172.

[7] T. Wisspeintner, T. van der Zant, L. Iocchi, and S. Schiffer, "RoboCup@Home: Scientific competition and benchmarking for domestic service robots," *Interaction Studies*, vol. 10, no. 3, pp. 393–428, 2009.

[8] D. Holz, A.-L. Jouen, M. Rajesh, J. Savage, K. Sugiura, L. Iocchi, J. R. del Solar, and T. van der Zant, "RoboCup@Home: Rules & regulations," http://purl.org/holz/rulebook.pdf, 2011.

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

[10] F. Hegger, C. Müller, Z. Jin, J. A. A. 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, Turkey, 2011.

[11] 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, Turkey, 2011.

[12] J. Bohren, R. Rusu, E. 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 the IEEE Int. Conf. on Robotics and Automation (ICRA)*, Shanghai, China, 2011, pp. 5568–5575.

[13] M. Beetz, U. Klank, I. Kresse, A. Maldonado, L. Mösenlechner, D. Pangercic, T. Rühr, and M. Tenorth, "Robotic roommates making pancakes," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, Bled, Slovenia, 2011, pp. 529–536.

[14] A. Jain and C. C. Kemp, "EL-E: an assistive mobile manipulator that autonomously fetches objects from flat surfaces," *Autonomous Robots*, vol. 28, no. 1, pp. 45–64, 2010.

[15] S. Srinivasa, D. Ferguson, C. Helfrich, D. Berenson, A. Collet, R. Diankov, G. Gallagher, G. Hollinger, J. Kuffner, and J. M. Vandeweghe, "HERB: A home exploring robotic butler," *Autonomous Robots*, vol. 28, no. 1, pp. 5–20, 2010.

[16] Z. Xue, S. Ruehl, A. Hermann, T. Kerscher, and R. Dillmann, "An autonomous ice-cream serving robot." in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, Shanghai, China, 2011, pp. 3451–3452.

[17] 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 the IEEE Int. Conf. on Robotics and Automation (ICRA)*, Shanghai, China, 2011, pp. 3443–3444.

[18] J. Stückler and S. Behnke, "Integrating indoor mobility, object manipulation, and intuitive interaction for domestic service tasks," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, Paris, France, 2009, pp. 506–513.

[19] 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 European Conf. on Mobile Robots (ECMR)*, Örebro, Sweden, 2011, pp. 177–182.

[20] S. Cousins, B. Gerkey, K. Conley, and Willow Garage, "Sharing software with ROS," *IEEE Robotics and Automation Magazine*, vol. 17, no. 2, pp. 12–14, 2010.

[21] D. Fox, "Adapting the sample size in particle filters through KLD-sampling," *Int. Journal of Robotics Research (IJRR)*, vol. 22, no. 12, pp. 985–1003, 2003.

[22] P. Hart, N. Nilson, and B. Raphael, "A formal basis for the heuristic determination of minimal cost paths," *IEEE Transactions of Systems, Science, and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

[23] G. Grisetti, C. Stachniss, and W. Burgard, “Improved techniques for grid mapping with Rao-Blackwellized particle filters,” *IEEE Transactions on Robotics*, vol. 23, no. 1, pp. 34–46, 2007.

[24] J. Stückler and S. Behnke, “Compliant task-space control with back-drivable servo actuators,” in *Proc. of the RoboCup International Symposium*, Istanbul, Turkey, 2011.

[25] G. Niemeyer and J.-J. E. Slotine, “A simple strategy for opening an unknown door,” in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, Albuquerque, NM, USA, 1997, pp. 1448–1453.

[26] 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. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, Anchorage, AK, USA, 2010, pp. 1807–1814.

[27] R. B. Rusu, N. Blodow, Z. C. Marton, and M. Beetz, “Close-range scene segmentation and reconstruction of 3D point cloud maps for mobile manipulation in human environments,” in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, St. Louis, MO, USA, 2009, pp. 1–6.

[28] D. Holz, S. Holzer, R. B. Rusu, and S. Behnke, “Real-time plane segmentation using RGB-D cameras,” in *Proceedings of the RoboCup International Symposium*, Istanbul, Turkey, July 2011.

[29] M. A. Fischler and R. C. Bolles, “Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography,” *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.

[30] C. Borst, M. Fischer, and G. Hirzinger, “Efficient and precise grasp planning for real world objects,” in *Multi-point Interaction with Real and Virtual Objects*, ser. Springer Tracts in Advanced Robotics, 2005, vol. 18, pp. 91–111.

[31] A. Miller and P. Allen, “GraspIt! A versatile simulator for robotic grasping,” *IEEE Robotics and Automation Magazine*, vol. 11, no. 4, pp. 110–122, 2004.

[32] K. Hsiao, S. Chitta, M. Ciocarlie, and E. G. Jones, “Contact-reactive grasping of objects with partial shape information,” in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Taipei, Taiwan, 2010, pp. 1228–1235.

[33] J. Stückler and S. Behnke, “Following human guidance to cooperatively carry a large object,” in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, Bled, Slovenia, 2011, pp. 218–223.

[34] J. Stückler and S. Behnke, “Improving people awareness of service robots by semantic scene knowledge,” in *Proc. of the RoboCup International Symposium*, Singapore, Singapore, 2010.

[35] D. Droeschel, J. Stückler, D. Holz, and S. Behnke, “Towards joint attention for a domestic service robot – person awareness and gesture recognition using time-of-flight cameras,” in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, Shanghai, China, 2011, pp. 1205–1210.