

Towards Robust Mobility, Flexible Object Manipulation, and Intuitive Multimodal Interaction for Domestic Service Robots

Jörg Stückler, David Droeschel, Kathrin Gräve, Dirk Holz,
Jochen Kläß, Michael Schreiber, Ricarda Steffens, and Sven Behnke

Rheinische Friedrich-Wilhelms-Universität Bonn
Computer Science Institute VI: Autonomous Intelligent Systems
Friedrich-Ebert-Allee 144, 53113 Bonn, Germany
{ stueckler | droeschel | graeve | holz | schreiber | behnke } @ ais.uni-bonn.de
{ klaehs | steffens } @ cs.uni-bonn.de
<http://www.NimbRo.net/@Home>

Abstract. In this paper, we detail the contributions of our team NimbRo to the RoboCup @Home league in 2011. We explain design and rationale of our domestic service robot Cosero that we used for the first time in a competition in 2011. We demonstrated novel capabilities in the league such as real-time table-top segmentation, flexible grasp planning, and real-time tracking of objects. We also describe our approaches to human-robot cooperative manipulation and 3D navigation. Finally, we report on the use of our approaches and the performance of our robots at RoboCup 2011.

1 Introduction

The RoboCup @Home league has been established in 2006 to research and benchmark autonomous service robots in everyday scenarios. In the first years of the league, basic capabilities of the robots have been tested. The robots had to show object recognition and grasping, safe indoor navigation, and basic human-robot interaction (HRI) skills such as speech recognition and synthesis. Progress in the league allowed to introduce more complex test procedures in 2010.

Our team NimbRo participates in the @Home league since 2009. In the first year, we could demonstrate basic mobile manipulation and HRI skills with our self-constructed robots Dynamaid and Robotinho. We could reach the third place in the competition and received the Innovation Award for innovative robot body design, empathic behaviors, and robot-robot cooperation. In the second year, we participated with a mechanically improved version of Dynamaid. We tackled the more complex tests and showed many new mobile manipulation and human-robot interaction skills such as gesture recognition. We also demonstrated as the first team the opening and closing of a refrigerator at RoboCup. Overall, we could reach the second place in 2010.




Fig. 1. The cognitive service robot *Cosero* cooperatively carries a table with a user and bakes omelett at the RoboCup@Home finals 2011 in Istanbul.

In this year’s competition, we participated with Dynamaid and its successor, Cosero. While Cosero still retains the light-weight design of Dynamaid, we improved its construction and appearance significantly and made it more precise and stronger actuated. 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 2011. In this paper, we summarize our main novel contributions to the RoboCup@Home league. We detail the construction of Cosero, and the algorithms we developed in the context of the league.

2 Design of Cognitive Service Robot Cosero

Our everyday environments are adapted to the specific capabilities and constraints of the human body. When robots perform similar tasks like humans in such environments, it is a natural design rationale to equip the robots with human-like motion and perception abilities. A further advantage of a human-like body is that the robot’s behavior is predictable and can easily be interpreted by humans. In everyday scenarios, robots also may interact physically with humans. This imposes requirements on the safety of such a robot. A light-weight design makes a household robot inherently less dangerous than a heavy-weight industrial-grade robot.

We focused the design of our robots Dynamaid and Cosero (s. Fig. 1) on such requirements. Cosero’s mobile base has a small footprint of 59×44 cm and drives omnidirectionally. This allows Cosero 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 by approx. 0.9 m. This allows the robot to manipulate on similar heights like humans – even on the floor.

Cosero has been constructed from light-weight aluminum parts. All joints in the robot 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. The robot’s main computer is a HP Pavillion dv6 notebook with an Intel i7-Q720 processor.

Cosero perceives its environment with a variety of complementary sensors. The robot senses the volume in front of it in 3D with a Microsoft Kinect RGB-D camera in its head that is attached to the torso with a pan-tilt unit in the neck. 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. One Hokuyo URG-04LX is placed on a roll actuator in the lower torso. Aligned horizontally, it detects objects on horizontal surfaces such as tables or shelves. In vertical alignment, it measures distance and height of these objects. A further Hokuyo URG-04LX measures objects in a height of ca. 4 cm above the floor. The main sensor for 2D localization and mapping is a SICK S300 sensor in the mobile base which perceives the environment in the horizontal plane at a height of approx. 27 cm. In the upper torso, we mounted a Hokuyo UTM-30LX laser-range scanner on a tilt actuator to acquire precise range measurements in up to 30 m for 3D obstacle avoidance, mapping, and localization. The laser-range scanners are also useful to track persons in the robot’s surroundings.

To improve the robustness of manipulation, the robot can measure the distance to obstacles directly from the grippers. We attached infrared distance sensors to each gripper that point downward and forward in the finger tips. Another sensor in the palm measures distance to objects within the gripper.

Finally, the sensor head also contains a shotgun microphone for speech recognition. By this, the robot points the microphone towards human users and at the same time directs its visual attention to the user. We attached a human face mask to support the interpretation of the robot’s gaze by the human.

3 Mobile Manipulation

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

3.1 Perception

Real-Time Table-Top Segmentation: In household environments, objects are usually located on planar surfaces such as tables. We therefore base our object detection pipeline on fast horizontal plane segmentation of the depth



Fig. 2. Object detection. Left: example table top setting. Center: raw point cloud from the Kinect with RGB information. Right: each detected object is marked with a distinct color.




Fig. 3. Grasp planning. Left: object shape properties. The arrows mark the principal axes of the object. Center: we rank feasible, collision-free grasps (red, size prop. to score) and select the most appropriate one (large, RGB-coded). Right: example grasps.

images of the Kinect [14]. Fig. 2 shows an exemplary result of our approach in a table-top scene. Our method 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 [9].

In order to process the depth images efficiently, we combine rapid normal estimation [7] with fast segmentation techniques. The normal estimation method exploits the principle of 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 [5] and find the objects by clustering the points above the convex hull of the support plane.

Grasp Planning: We investigate grasp planning to enlarge the set of graspable objects and to allow for obstructions by obstacles [14]. 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




Fig. 4. Learning of object models. Left: during training the user selects points (red dots) to form a convex hull around the object. Center: color and shape distribution modeled at 5 cm resolution. Lines indicate surface normals (color-coded by orientation). Right: color and shape distribution modeled at 2.5 cm resolution.

approach also may yield stable grasps when an object violates these assumptions. Fig. 3 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 and evaluate the grasps for feasibility under kinematic and collision constraints. The remaining grasps are ranked according to efficiency and robustness criteria. The best grasp is selected and finally executed with a parametrized motion primitive. For collision detection, we take a conservative but efficient approach that checks simplified geometric constraints.

Real-Time Object Tracking: When a robot interacts with objects, it has to estimate its pose with respect to the objects. Frequently, localization 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 change. The robot must then be able to detect the object in its current sensor view and estimate the relative pose of the object.

We develop methods for real-time tracking of objects with RGB-D cameras [12]. In our approach, we train a multi-resolution surfel map of the object (s. Fig. 4). The map is represented in an octree where each node stores a normal distribution of the volume it represents. In addition to shape information, we also model the color distribution in each node.

For fast object teach-in, we use checkerboard patterns laid out around the object. In the images, the user selects points on a convex hull around the object in the common plane of the checkerboards. The visual markers yield a precise map reference frame, in which various views on the object can be merged.




Fig. 5. Left: panorama image of an office. Center: 3D surfel map learned with our approach (surfel orientation coded by color). Right: 2D navigation map extracted from the 3D surfel map (gray: unknown, white: traversable, black: untraversable).

Once the map has been obtained, we build multi-resolution surfel maps with color information from new RGB-D images [13]. Then, we register this map to the object map 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-ups. We optimize the matching likelihood to find the most likely pose. Our efficient implementation supports real-time registration of RGB-D images on the robot’s main computer.

3D Perception for Navigation: Reliable obstacle avoidance is a prerequisite for safe navigation of a robot in indoor environments. We implemented means to incorporate 3D measurements from laser scanners and depth cameras [3]. We maintain the measurements in a 3D point cloud in the close surrounding of up to 10m. In order to handle dynamic objects, a temporal unlearning strategy discards measurements that are older than a fixed amount of time. When a sensor sweeps a map region with its field-of-view, the points in this volume are removed after shorter duration.

Many approaches to indoor localization and mapping use 2D laser scanners to acquire 2D footprints of the environment. Occupancy grid maps are used to represent the map, because they provide dense information about free and occupied space for localization and path planning. One problem of such 2D maps occurs in path planning, if untraversable obstacles cannot be perceived on the laser scanners 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 address these problems by building 3D maps of the environment [8]. Fig. 5 demonstrates example maps generated with our approach. We choose to represent the map in a 3D surfel grid. The robot acquires full 3D scans at several locations in the environment. Given the location w.r.t. the environment, we perform mapping in known poses to generate the surfel map from several scans. We obtain the trajectory estimate using a 2D SLAM approach (GMapping, [6]) and refine it by registering the 3D scans with ICP [1].




Fig. 6. Localization in dynamic environments. Left: exemplary result of the tracking performance of 2D localization (2D Grid) and our 3D localization approach (Surfel). Center: localization accuracy for pose tracking. Right: global localization accuracy.

Once the map has been obtained, we extract a 2D navigation map. First, we find the traversable cells in the 2D map by region growing in the 3D map with the robot’s scan poses as seed points. The region growing algorithm expands to cells when the gap between the surfels at the cell borders is small enough to traverse it. Additionally, we check if all cells are free within the robot’s height range.

For localization, we developed a 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. In both cases, we use a nearest-neighbor look-up grid to efficiently associate measurements with surfels.

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 concentrate the localization on planar structures, as they more likely occur in static environment parts.

Fig. 6 shows experimental results for pose tracking and global localization. In the experiments, we compare the reliability of standard Monte Carlo localization in 2D occupancy maps and our localization method which prefers measurements above the persons’ heights. Eight persons were randomly walking in the test environment. We quantify, how often and how accurate the localization methods estimate the final position of the trajectory. We use the SICK S300 laser scanner on the mobile base for 2D localization. For 3D localization, the laser scanner in the chest is continuously tilted. When the robot stands during a full sweep, the complete 3D scan is integrated. Otherwise, we use the immediate 2D scans. In pose tracking, we initialize the localization at the correct position. It can be seen that our approach localizes the robot more accurately. For global localization, we initialize the methods with a uniform distribution of 5000 particles. We evaluated global localization at 45 starting points in various trajectories. Global localization in the 2D map only succeeds in ca. 30% of the runs, whereas our approach achieves 97.5% success rate at a distance threshold of 0.5 m. While our approach yields superior results, it still retains the efficiency of 2D localization.

3.2 Behavior Control

Motion Control: The design of the mobile bases of our robots supports omnidirectional driving [9]. The linear and angular driving velocities can be set independently and can be changed to speeds within a continuous range. The drive consists of four differential drives each located at the corners of the rectangular base. We determine their steering direction and the individual wheel velocities from an analytical solution to the drive’s inverse kinematics.

The anthropomorphic arms support control in Cartesian coordinates. For this, we implemented differential inverse kinematics with redundancy resolution [9]. We also developed compliance control for the arms [11]. For our method, we exploit that the servo actuators are back-drivable and that the torque which the servo applies for position control can be limited. Compliance can be set for each direction in task-space separately. For example, the end-effector can be kept loose in both lateral directions while it keeps the other directions at their targets.

We implemented several motion primitives like grasping with one or two arms, pointing or waving gestures, and object placement. Motion primitives such as side- and top-grasps or pointing gestures can be parametrized in the target.

Mobile Manipulation Control: For mobile manipulation, we developed controllers that combine control of the drive and the arms with perception capabilities. Cosero can grasp objects on horizontal surfaces on tables and in shelves in a height range from ca. 0.3m to 1m [9]. It also carries the object and hands it to human users. We further 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. When handing an object over, the arms are compliant in upward direction so that the human can pull the object, the arm complies, and the object is released.

The robots can also open and close doors, if the door leaf can be moved without the handling of an unlocking mechanism. For example, fridges or cabinets are commonly equipped with magnetically locked doors that can be pulled open without special manipulation of the handle. To open a door, the robot drives in front of the door, detects the door handle with its 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 grasp the handle and moves forward while it holds the handle at its relative initial grasping pose. The arm deviates from this pose by the constraining motion of the door leaf, and the robot drives 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.

4 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.

4.1 Person Awareness

A key prerequisite for human-robot interaction is the robot’s awareness of the persons that surround it. Our robots maintain a belief on the location and identity of persons [10]. We implemented a multi-hypothesis tracker that initializes new person beliefs, when faces or upper bodies are detected at reasonable locations in the environment. Using laser scanners, the position and moving speed of the persons is then tracked at high frame rates.

4.2 Speech Recognition and Synthesis

Speech is the primary modality for the communication of complex statements between humans. We therefore support speech in our robots employing the Loquendo SDK. Its speech recognition is speaker-independent and uses a small-vocabulary grammar which we change with the dialog state. The grammar definition of the Loquendo speech recognition system allows to tag rules with semantic attributes. When speech is recognized, a semantic parse tree is provided that we process further. We use the parsed semantics to interpret sentences for complex commands and to generate appropriate behavior.

4.3 Gesture Recognition and Synthesis

An important non-verbal communication cue is the recognition and performance of gestures. We equipped our robots with several gestures. For example, the robot can draw a user’s attention to certain locations in the environment by simply pointing at it. Our robots can also perceive gestures such as pointing, showing of objects, or stop gestures [4]. The robots sense these gestures using the RGB-D camera. For pointing gestures, we accurately estimate the pointing direction from body features such as the position of the head, hand, shoulder, and elbow. We also investigated the use of Gaussian Process regression to learn an interpretation of the pointing direction [2] using the body features.

4.4 Human-Robot Cooperative Manipulation

We study physical interaction between a human user and a robot in a cooperative manipulation task [12]. 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 human actions.

In order to approach the table, the robot tracks its 6D pose with our real-time object tracking approach. The robot waits then until the human lifts the table. As soon as the human lifts the table, the robot measures a significant pitch of the table. Then, the robot also lifts the table and begins to hold the table compliant in the horizontal plane. The human user pulls and pushes the table into the desired directions, and the robot compensates for the displacement of its end-effectors by driving accordingly with its mobile base. When the target location of the table is reached, the user can simply put the table down which is detected by the robot.

5 Experiments 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 such that we achieved the first place in the competition.

In the *Robot Inspection and Poster Session* test in Stage I, Cosero and Dynamaid registered themselves. Meanwhile, we presented our work to leaders of other teams in a poster session. Overall, we have been awarded 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 showed that it can distinguish the 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 demonstrations. In this 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 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.

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. 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 about the location of the object. After successful execution, it worked on a task with erroneous information. It detected that the ordered object is not at the specified location, and went back to the user to report 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 teaboxes from a table. The objects could be chosen and placed by a jury member. The technical committee awarded us the

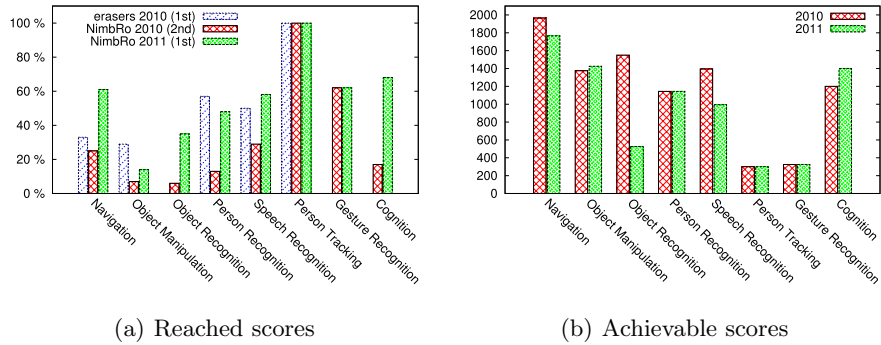


Fig. 7. Reached (a) and achievable (b) scores in the predefined test procedures per functionality.

highest score. We reached the finals with 8,462 points followed by WrightEagles from China with 6,625 points.

In the finals, we demonstrated the cooperative carrying of a table by Cosero and a human user. Then, a user showed Cosero where it finds a bottle of dough to make an omelett. Our robot then went to the cooking plate to switch it on. It succeeded partially in turning the plate on. Then, Cosero 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 our real-time object tracking method in order to approach the cooking plate. Meanwhile, Dynamaid opened a refrigerator, grasped a bottle of orange juice out of it, and placed the bottle on the breakfast table. Our performance received the best score by the jury consisting of members of the executive committee and external judges from science and the media.

Fig. 7 summarizes the scores achieved for individual functionalities as proposed in [15]. 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 improved most functionalities compared to 2010 and achieved well in developing a balanced domestic service robot system.

6 Conclusion

In this paper, we presented our developments for the RoboCup@Home league in 2011. We detailed our approaches to real-time scene segmentation, object tracking, 3D navigation, and human-robot cooperative manipulation. We use the RoboCup@Home competitions to evaluate our methods in a realistic setting. With our domestic service robots, we won the competitions in 2011.

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 of the robot 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 could devise new methods to grasp smaller objects or to use smaller tools.

References

1. P. J. Besl and N. D. McKay. A method for registration of 3-D shapes. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 14:239–256, 1992.
2. D. Droeschel, J. Stückler, and S. Behnke. Learning to interpret pointing gestures with a time-of-flight camera. In *Proc. of the 6th ACM/IEEE Int. Conf. on Human-Robot Interaction (HRI)*, 2011.
3. D. Droeschel, J. Stückler, D. Holz, and S. Behnke. Using time-of-flight cameras with active gaze control for 3D collision avoidance. In *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2010.
4. 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)*, 2011.
5. 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*, 24(6):381–395, 1981.
6. G. Grisetti, C. Stachniss, and W. Burgard. Improved techniques for grid mapping with Rao-Blackwellized particle filters. *IEEE Trans. on Robotics*, 23(1), 2007.
7. 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.
8. 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.
9. 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.
10. J. Stückler and S. Behnke. Improving people awareness of service robots by semantic scene knowledge. In *Proc. of the 14th Int. RoboCup Symposium*, 2010.
11. J. Stückler and S. Behnke. Compliant task-space control with back-drivable servo actuators. In *Proc. of the 15th RoboCup International Symposium*, 2011.
12. J. Stückler and S. Behnke. Following human guidance to cooperatively carry a large object. In *Proc. of the IEEE-RAS Int. Conf. on Humanoid Robots*, 2011.
13. J. Stückler and S. Behnke. Robust real-time registration of RGB-D images using multi-resolution surfel maps. In *Proc. of 7th German Conf. on Robotics (ROBOTIK)*, 2012.
14. 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 5th European Conference on Mobile Robots (EMCR)*, 2011.
15. 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.