

homer@UniKoblenz

Viktor Seib, Nicolai Wojke, Raphael Memmesheimer, Daniel Müller, Sebastian Stümper,
Florian Sattler, Alexander Kreutz, Sebastian Brockmann, Alruna Veith, Arne Peters,
Denise Dünnebier, Eva Kreckel, Ivanna Mykhalchyshyna, Dietrich Paulus

Active Vision Group (AGAS)
University of Koblenz and Landau
Universitätsstr. 1, 56070 Koblenz, Germany
vseib@uni-koblenz.de, <http://homer.uni-koblenz.de>

Abstract—This paper describes the robot hardware and software used by team homer@UniKoblenz of the University of Koblenz and Landau, Germany, for the participation at the RoCKIn@Home competition 2014 in Toulouse. A special focus is put on detailed hardware and software description and novel scientific achievements. For navigation and mapping we use well-established SLAM techniques based on particle filter and grid maps. Object recognition is achieved by clustering of local invariant features. Three-dimensional scans of the environment are acquired using a RGB-D sensor mounted on a pan-tilt unit. In order to manipulate objects a 6 DOF industrial-grade robotic arm is employed. An abstraction layer in our software architecture allows for task planning and general purpose task execution. For the improvement of human-robot interaction we developed a generic face model that is synchronized to speech and can show seven different face expressions. This robot face is available as a ROS-node for other teams. Some of the novelties of this year is a new robotic platform and a complete change to ROS that now replaces our previously used software architecture. Further, our object recognition algorithm was published as ROS package. Furthermore, improvements were made in speech filtering, people detection and tracking as well as face expression recognition.

I. INTRODUCTION

Our team homer@UniKoblenz was formed to participate in the RoboCup@Home competition. Our team already participated successfully as finalist in Suzhou, China (2008), Graz, Austria (2009) in Singapur (2010), where it was honored with the RoboCup@Home Innovation Award, in Mexico-City, Mexico (2012), where it was awarded the RoboCup@Home Technical Challenge Award and in Eindhoven (2013). Further, we participated in stage 2 at the RoboCup@Home World Championship in Istanbul, Turkey (2011).

Besides RoboCup@Home we competed in the RoboCup Rescue league with our robot *Robbie*, where our team won the Interleague Mapping Challenge award (2010) and became German Champion in Rescue Autonomy for the 5th time in 2011.

A part of our current team participated in the RoCKIn@Home camp 2014 where we won the best final presentation award.

In 2014 we plan to attend the RoCKIn@Home competition in Toulouse with our robot *Lisa* (Fig. 1). Our team will be presented in the next section. Section III is devoted to the hardware used for Lisa. In Section IV our software architecture, autonomous navigation and human-robot interaction will be described. Finally, Section V will summarize this paper.



Fig. 1. Current setup of our robot Lisa.

II. ABOUT OUR TEAM

The members of team homer@UniKoblenz are mostly undergraduate and Master students from the University of Koblenz and Landau, Germany. They develop and improve our robot in practical courses offered by the Active Vision Group (AGAS). Head of the group is Prof. Dr.-Ing. Dietrich Paulus.

Since most practical courses last one semester only, it is natural that most of the team members are replaced on a regular basis. Each practical course is supervised by PhD students who act as scientific supervisors, but also incorporate their research into the project.

A. Team Members and their Contributions

This year, team *homer@UniKoblenz* consists of the following members, listed with their contributions:

Viktor Seib: team leader, scientific supervisor, object recognition

Nicolai Wojke: scientific supervisor, object tracking

Raphael Memmesheimer: system integration, 3D mapping

Daniel Müller: speech recognition

Sebastian Stümper: person tracking

Florian Sattler: hardware integration

Alexander Kreutz: hardware integration

Sebastian Brockmann: grasp planning

Arne Peters: grasp planning

Alruna Veith: people detection

Denise Dünnebie: graphical user interface

Eva Kreckel: gesture recognition

Ivanna Mykhalchyshyna: grid mapping and navigation

B. Focus of Research

Our interests in research are sensor fusion for people tracking and people detection, object recognition and categorization, face and gesture recognition as well as object manipulation using a robotic arm.

III. HARDWARE

A. Robot Platform

Currently, we are migrating to a new robotic platform, the CU-2WD-Center, manufactured by UlrichC¹. It is equipped with odometry sensors and has a 2 wheel drive in contrast to the 4 wheel drive of our previous platform. It allows the robot to turn on the spot with significantly less friction and thus, saves battery power and preserves the motors. On top of the platform, we have installed a prototype framework, which was designed and built by the Center of Excellence of the Chamber of Crafts in Koblenz and is able to carry additional sensors and a notebook running the control software.

B. Sensors and additional Hardware

1) *Notebook*: The software of the robot runs on a Lenovo W520 notebook equipped with an Intel Core i7-2670QM processor and 12 GB of RAM using Ubuntu Linux 13.04 as operating system.

2) *SICK LMS100 laser range finder*: The SICK LMS100 is mounted at the bottom and generates 180° scans. It has an adjustable angular resolution, while its maximal measured distance is 20 m. It is used for mapping, localization and people tracking.

3) *DirectedPerception PTU-D46 pan-tilt unit*: The DirectedPerception PTU-D46 is mounted on top of the robot's neck. It is able to rotate 159° in each direction and to tilt from +31° to -47° out of a horizontal position. The angular resolution is 0.012857°. Further sensors are attached on top of this unit.

4) *Rode VideoMic Pro Microphone*: The microphone Rode VideoMic Pro is used for speech input. It is a light-weight microphone and is mounted on the pan-tilt unit of our robot. We use the Jack audio server² to capture speech and pass it either directly to the system or a noise filter. For filtering, the Quantile-based noise estimator [1] is tested. This audio setup and speech filtering was recommended by and tested in cooperation with team Golem from the UNAM, Mexico, that participates in the RoboCup@Home.

5) *Canon PowerShot SX100 IS*: In earlier experiments we have found out that because of their low resolution the color images provided by the Kinect are not suitable for reliable object recognition. Thus, this off-the-shelf digital camera is mounted on top of the pan-tilt unit to capture high resolution (4 to 8 megapixels) images for object recognition.

6) *Microsoft Kinect*: The Microsoft Kinect is attached to the pan-tilt unit and provides depth and color images of size 640x480 pixels at 30Hz. The depth sensor operation range lies between 0.8 m and 3.5 m at a horizontal field of view of 58° and a vertical field of view of 45°. It also has two microphones and a tilt motor for sensor adjustment. However, the robot is equipped with another microphone for speech recognition. Its tilt motor is not used because of the pan-tilt unit that the Kinect is mounted on.

7) *Neuronics Katana 400HD*: The Katana 400HD is a 6 DOF industrial-grade robot arm. It is attached to our robot's body plate and used to manipulate objects on tables and other furniture of similar height. With an accuracy of 1 mm and a length of 90 cm, it enables us to perform delicate manipulation tasks on light-weight objects.

The end effector is a custom setup and consists of 4 Festo Finray-fingers. It is safe for interaction with humans.

IV. TECHNOLOGY AND SCIENTIFIC CONTRIBUTION

A. System Core Architecture

In the last years we used a hybrid software architecture that consisted of a self developed architecture Robbie [2] and an interface to ROS³ (Figure 2). For the first time this year, we completely migrated all of our software to ROS to benefit from a homogeneous architecture. Further, using this standard middleware allows us to use code already available for ROS, but also to share some of our components with the community.

B. Graphical Interface

We also migrated our GUI to ROS. It can be run directly on the robot or on a different computer via WLAN. The user interface is realized using Qt4 and OpenGL, further, with the migration to ROS we included the rviz plugin. A visualization example is given in Fig. 3.

C. Simultaneous Localization and Mapping

To enable users without technical knowledge to use the robot and to ease the setup procedure, it has to be able to create a map without the help of a human. For this purpose, our robot continuously generates and updates a 2D map of its

¹Manufacturer of our new robotic platform: <http://www.ulrichc.de>

²Jack audio server: <http://jackaudio.org/>

³Documentation of ROS <http://www.ros.org/>

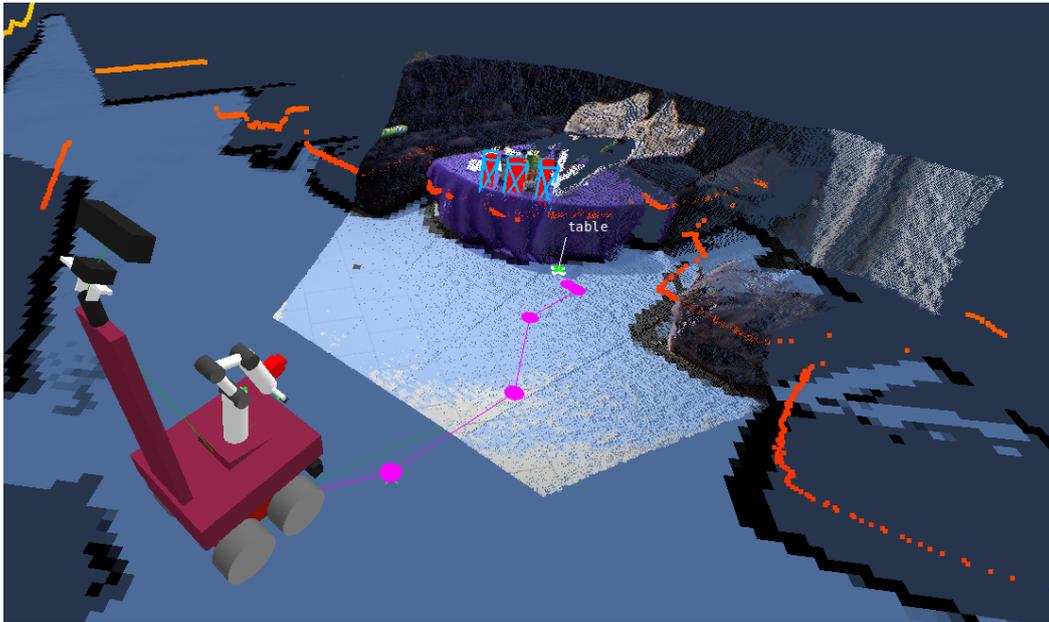


Fig. 3. Visualization of RGBD and laser sensor data. Further, a planned path to the point of interest “table” as well as grippable objects on the table are visualized.

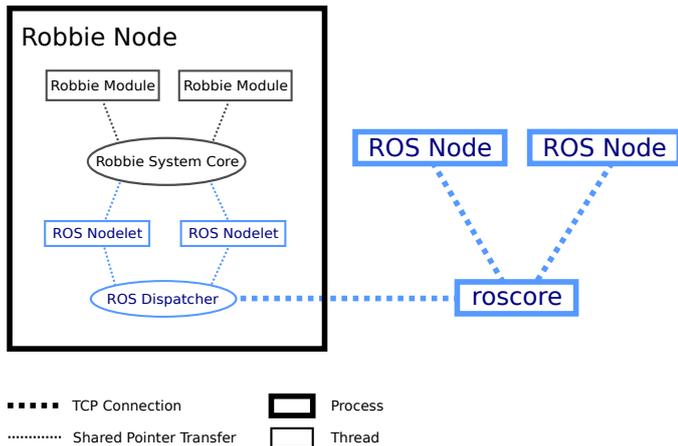


Fig. 2. Schematic description of the communication between Robbie and ROS.

environment based on laser scans. Figure 4 shows an example of such a map.

To detect obstacles below and above the LRF-plane we use Kinect data to augment the occupancy map.

We plan to publish our mapping and navigation algorithms as ROS packages in the second half of 2014.

D. Navigation in Dynamic Environments

In real-life situations, a static grid map is not sufficient for navigating through an everyday environment, as due to the movement of persons and other dynamic obstacles, an occupancy map that only changes slowly in time does not provide sufficient information.

Thus, our navigation system, which is based on Zelinsky’s path transform [3], [4], always merges the current laser range



Fig. 4. Real-time maps of the RoboCup 2008 (left) and 2009 (right) @Home arena.

scan as a frontier into the occupancy map. A once calculated path is then checked against obstacles in small intervals during navigation, which can be done at very little computational expense. If an object blocks the path for a given interval, the path is re-calculated. This approach allows the robot to efficiently navigate and avoid obstacles in highly dynamic environments.

E. Human-Robot Interface

The robot is equipped with speakers and a microphone which enables communication via speech interface. In addition, it has a small screen that is used to display facial expressions and state information. On the software side, we decided to use two open source libraries: *pocketsphinx*⁴ for speech recognition and *festival*⁵ for speech synthesis.

⁴Speech recognition system *pocketsphinx* <http://www.speech.cs.cmu.edu/pocketsphinx/>

⁵Speech synthesis system *festival* <http://www.cstr.ed.ac.uk/projects/festival/>

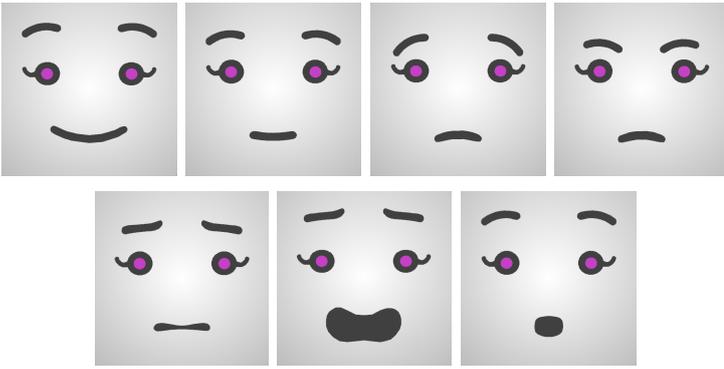


Fig. 5. Animated face of our service robot Lisa. The depicted face expressions are (from left to right): happy, neutral, sad, angry, disgusted, frightened, and surprised.

We have designed a concept of a talking robot face that is synchronized to speech via mouth movements. For this feature we are using the open source library Ogre3D⁶ for visualisation. Furthermore, the face is modelled with Blender⁷ and exported via the Ogre Mesh Exporter⁸ for the use with Ogre.

To include the robot face into our software, we have created a ROS node for this application. We extended the robot face to show seven different face expressions (Figure 5). Further, we provide two similar face meshes, a female and a male one. The colors and the voice (female or male) can be configured via a configuration file without recompiling the application. The robot face has been released as an open source package for ROS⁹. Thus, it is possible to apply it for any robot. Furthermore, we conducted a broad user study to test how people perceive the shown emotions. The results can be found in [5].

To further enhance communication with humans a robot has to detect their emotions. One step towards this goal is the recognition of facial expressions. We use an approach based on the facial action coding system and action units. In the first step features are extracted from the detected face and classified into action units. The subsequent step analyzes the obtained action units and assigns a facial expression. Both classification steps use a support vector machine. Our system is able to recognize the same 7 face expressions that Lisa can exhibit. This approach works without prior calibration with a neutral face expression of the detected person.

F. Object detection

Objects for mobile manipulation are detected by first segmenting horizontal planes as table hypotheses. Subsequently, all points above the plane are clustered and the resulting clusters considered as objects.

Transparent objects (in our case drinking glasses) are detected by making use of one fault of the Kinect sensor. The structured light emitted by the Kinect is scattered in transparent objects providing no valid depth information. We segment

areas with no depth data and compare them to holes in detected planes to extract contours and match them with drinking glass contour templates. Since the supporting table plane around the transparent objects has valid depth information, a size and location estimation of the transparent objects is obtained and used for grasping.

G. Object and Face Recognition

The object recognition algorithm we use is based on Speeded Up Robust Features (SURF) [6], which are local scale-invariant features of gray images. First, features are matched between the trained image and the current camera image based on their euclidean distance. A threshold on the ratio of the two nearest neighbours is used to filter unlikely matches. Then, matches are clustered in hough space using a four dimensional histogram using their position, scale and rotation. This way, sets of consistent matches are obtained. The result is further optimized by calculating a homography between the matched images and discarding outliers. Our system was evaluated in [7] and shown as suitable for fast training and robust object recognition.

With this object recognition approach we won the Technical Challenge 2012. Our object recognition is available as open source software in ROS¹⁰. A detailed description of the approach is provided on the website and in [8] as well.

H. People Detection

People are detected by the combination of three sensors. The laser range finder is used to detect legs. The RGB image of the Kinect camera provides data for face detection. We use the face detection algorithm implemented in the OpenCV library. Finally, the Kinect depth camera allows to detect silhouettes of persons. For a person to be detected, not every sensor has to see the person. However, the more sensors see a person the higher the probability to really encounter a person at the position in question.

The silhouette detection works in three steps. First, candidate silhouettes are found. In order to do this, the depth data is split in overlapping slices along the forward axis. Each slice is then tested for point clusters roughly matching a human silhouette (size and round shape of the head). In a second step low-frequency fourier features are extracted from each candidate silhouette. In the last step a support vector machine classifies the feature vector as a non-human object, a lateral silhouette or as frontal silhouette.

I. People Tracking

For the perception of people we combine the three methods described in IV-H. In case that these sensors yield similar points, the output is merged. Thus, it is assumed that the points belong to the same person. The obtained person positions are used as input for the particle filter which estimates the movement and position of the people.

⁶Open-source graphics rendering engine Ogre3D <http://www.ogre3d.org/>

⁷Free open source 3D content creation suite Blender <http://www.blender.org/>

⁸Ogre Mesh Exporter <http://www.ogre3d.org/tikiwiki/Blender+Exporter>

⁹Robot Face: <http://www.ros.org/wiki/agas-ros-pkg>

¹⁰Object Recognition software and detailed description: <http://www.ros.org/wiki/agas-ros-pkg>

J. Object manipulation

Our system detects planes in the acquired 3D point cloud. This information is used to find euclidean point clusters on top of planes which fit into the gripper and are thus regarded as candidates for grasping.

If a specific object has to be grasped, one of the color cameras is adjusted to face the object cluster. The object detection algorithm is executed on the region of interest defined by the object cluster's bounding box.

The movement planning for our robotic arm is performed using an approach operating directly in working space. Chaining motion primitives, our path planner builds a graph from the starting position to the goal. The planning can be optimized towards specific objectives. These are performing a smooth path or keeping a maximum distance from obstacles using heuristic and cost functions [9].

V. CONCLUSION

In this paper, we have given an overview of the approach used by team homer@UniKoblenz for the RoCKIn@Home competition. We presented a combination of out-of-the box hardware and sensors and a custom-built robot framework. We explained the use of well-established techniques like SLAM based on a particle filter and a grid map. Further, we explained our approach for object recognition using matching and clustering of local invariant features and the ability to detect and manipulate objects with a 6 DOF robotic arm. Finally, we presented our software contribution to the community. So far, this is the animated robot face and the object recognition algorithm based on SURF. Both packages can be downloaded at <http://wiki.ros.org/agas-ros-pkg>.

REFERENCES

- [1] V. Stahl, A. Fischer, and R. Bippus, "Quantile based noise estimation for spectral subtraction and wiener filtering," in *Acoustics, Speech, and Signal Processing, 2000. ICASSP'00. Proceedings. 2000 IEEE International Conference on*, vol. 3. IEEE, 2000, pp. 1875–1878.
- [2] S. Thierfelder, V. Seib, D. Lang, M. Häselich, J. Pellenz, and D. Paulus, "Robbie: A message-based robot architecture for autonomous mobile systems," *INFORMATIK 2011-Informatik schafft Communities*, 2011.
- [3] A. Zelinsky, "Robot navigation with learning," *Australian Computer Journal*, vol. 20, no. 2, pp. 85–93, 5 1988.
- [4] A. Zelinsky, "Environment exploration and path planning algorithms for a mobile robot using sonar," Ph.D. dissertation, Wollongong University, Australia, 1991.
- [5] V. Seib, J. Giesen, D. Grüntjens, and D. Paulus, "Enhancing human-robot interaction by a robot face with facial expressions and synchronized lip movements," in *21st International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision*, V. Skala, Ed., 2013. [Online]. Available: <http://www.uni-koblenz.de/~agas/Documents/Seib2013EHI.pdf>
- [6] H. Bay, T. Tuytelaars, and L. Van Gool, "SURF: Speeded up robust features," *ECCV*, pp. 404–417, 2006. [Online]. Available: <http://www.vision.ee.ethz.ch/~surf/eccv06.pdf>
- [7] P. Decker, S. Thierfelder, D. Paulus, and M. Grzegorzec, "Dense Statistic Versus Sparse Feature-Based Approach for 3D Object Recognition," in *10th International Conference on Pattern Recognition and Image Analysis: New Information Technologies*, vol. 1. Moscow: Springer MAIK Nauka/Interperiodica, 12 2010, pp. 181–184.
- [8] V. Seib, M. Kusenbach, S. Thierfelder, and D. Paulus, "Object recognition using hough-transform clustering of surf features (report on the technical challenge of the robocup@home 2012, mexico-city)," 2013. [Online]. Available: <http://www.uni-koblenz.de/~agas/Documents/Seib2013ORU.pdf>
- [9] B. Cohen, S. Chitta, and M. Likhachev, "Search-based planning for manipulation with motion primitives," Anchorage, Alaska, 2010.