

RoboCup 2010 - homer@UniKoblenz (Germany)

David Gossow, Nicolai Wojke, René Bing, Urs Buchholz, Robin Schrage, Andreas Mützel, Kevin Read, Susanne Thierfelder, Sebastian Vetter, Dietrich Paulus

Active Vision Group
University of Koblenz-Landau
Universitätsstr. 1
56070 Koblenz, Germany
dgossow@uni-koblenz.de
<http://www.uni-koblenz.de/~robbie>

Abstract. This paper gives a description of the robot hardware and software used by team homer@UniKoblenz, which participates in RoboCup@home 2010. A Special Focus is put on novel scientific achievements and newly developed features with respect to last year's competition. For navigation and mapping we use well-established SLAM techniques based on particle filters and grid maps. Object recognition is achieved by clustering of local invariant features. Three-dimensional scans of the environment are acquired using a combination of a two-dimensional laser range finder attached to a pan-tilt unit as well as a time-of flight camera. To manipulate objects a 2-DOF gripper and a 6-DOF industrial-grade robotic arm are employed. Path planning and collision avoidance for the robotic arm is done using a novel graph-based search algorithm based on motion primitives in working space. Finally, we use the time-of-flight camera's depth images to perform gesture recognition.

Introduction

homer@UniKoblenz is a team of researchers and students from the University of Koblenz-Landau, Germany. We are aiming to participate in the RoboCup@home league. Our mobile system called "Lisa" was developed during the last five years. In this period, eight practical courses were held in which the robot hardware was assembled and the software developed, mainly by students. The robot was first used to guide persons to various places within the university, then as a transport vehicle. Finally, it started participating in the RoboCup Rescue league and, since RoboCup German Open 2008, in the @home league.

1 About our team

1.1 Team members and their contributions

The members of team *homer@UniKoblenz* are

- David Gossow: team leader
- Nicolai Wojke: technical chief designer
- René Bing: public relations
- Robin Schrage: quality assurance
- Urs Buchholz: hardware assistant
- Andreas Mützel: gesture recognition
- Kevin Read: robotic arm integration
- Susanne Thierfelder: face and object recognition

The @home team is supported by the Centre of Excellence of the Chamber of Crafts in Koblenz¹ (see section 2.1).

¹ <http://www.hwk-kompetenzzentrum.de>



Fig. 1. Our robot “Lisa”

1.2 Focus of research

Our interests in research are grid-based 2D mapping, localization and navigation, system architecture for automobile systems, real-time visualization of sensor data and system states, laser-based person tracking, object, face and gesture recognition as well as object manipulation using a manipulator.

2 Hardware

2.1 Robot platform

Pioneer3-AT As platform an MobileRobots Pioneer3-AT ² is used. It provides sonar and odometry sensors as described below. It is equipped with four air-filled tires having a diameter of 21.5 cm which can be controlled individually, allowing the robot to turn on the spot while maintaining a high degree of stability. Attached to the platform is a 2-DOF gripper, which will be used to pick up objects from the floor.

Home Robot Framework On top of the P3-AT, we have installed a prototype framework, which was designed and built by the Centre of Excellence of the Chamber of Crafts in Koblenz and will carry additional sensors and a notebook running the control software. It is designed in a way that local crafts enterprises can manufacture it using modern industrial materials like aluminium and different composites.

2.2 Sensors and additional Hardware

Notebook The software runs on a LG P310 Camini 8400 notebook equipped with a 2,26 Ghz Intel Processor and 3 GB of RAM using Ubuntu Linux as operating system.

² <http://www.activerobots.com>

Odometry sensors The Pioneer robot has a built in system to report the robot pose, based on the readings from the wheel encoders. We use this data as a rough approximation of the relative robot motion between two SLAM iterations, which is then refined by laser scan matching and a particle filter.

Sonar sensors Our Pioneer 3 AT robot platform has two sonar rings (one scanning the front and one scanning the back) with eight sonar sensors each. In total they cover 360 degrees with low resolution and larger gaps in between. The sonar sensors are used for collision avoidance and detection of transparent obstacles during the mapping and localization process, which are not discovered by the laser range finders (LRF).

SICK LMS100 laser range finder (New in 2010) The SICK LMS100 is mounted at the bottom and generates 270 degree scans. It has an adjustable angular resolution, while its maximal measured distance is 20 m. It is used for mapping, localization and people tracking.

Hokuyo URG04-LX laser range finder The Hokuyo laser range finder generates 240 degree scans that measure the distance of the closest objects. It has an angular resolution of 0.36° and a maximal measured distance of 5.6 m. The Hokuyo URG04-LX is mounted on the pan-tilt unit next to the camera such that 3d scans can be generated by step-wise tilting the unit. It is further used in combination with the second LRF to track people (see section 3.8).

DirectedPerception PTU-D46 pan-tilt unit The DirectedPerception PTU-D46 is mounted on the top of the robot's neck. It is able to rotate 159 degrees in each direction and to tilt from $+31^\circ$ to -47° out of a horizontal position. The angular resolution is 0.012857° . Sensors are attached on top of this unit.

Philips 1300NC and 1330NC color cameras Two color cameras with 1.3 and 2.0 megapixels are mounted on the robot for object and face recognition purposes. The Philips 1300NC is connected to the robot platform and used to capture images of objects in reach of the 2-DOF gripper. The 1330NC is connected to the pan-tilt unit. It is used to learn and recognize people.

Neuronics Katana 400HD (New in 2010) The Katana 400 is a 6-DOF industrial-grade arm robot. 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 concerning light-weight objects. The end effector is a standard pincher gripper and is safe for interaction with humans.

MESA Imaging SwissRanger 4000 (New in 2010) The MESA Imaging SwissRanger 4000 time-of-flight camera is attached to the pan-tilt unit and used for gesture recognition. It provides depth images of 174×144 pixels with a framerate of at maximum 50 fps.

3 Technology and Scientific Contribution

3.1 System Core Architecture

Our software is built upon a specially developed generic core, which is independent from hardware and purpose of the system. The core's main task is the exchange of messages in the system.

There are modules, each running as a single thread, which can be connected to the core and subscribe to messages. To keep these modules independent from each other the communication between them requires the core as a node. This means a message is sent to the core by a module and the core forwards the message to all the modules, which subscribed to this type of message.

Modules act as glue code, linking so-called workers and devices to the core system (see Fig. 2). They are supposed to handle messages and distribute data, providing little additional functionality. A Worker is a small set of program code (usually an object class) which mainly provides computing functionality, for example implementing a certain algorithm. In contrast, devices mainly provide access to a hardware components.

The system is highly configurable via a central registry stored as XML file. The registry contains various profiles, each one specifying its own set of modules to be loaded and configuration parameters to use.




Fig. 2. The Architecture of the Robbie system.

3.2 Graphical Interface




Fig. 3. The user interface of the operator station showing various visualizations of sensor data and system states. The application can be used for real-time surveillance of the robot as well as for playing back recorded sensor log files. In the @home league, this is mainly used by the developers for testing and evaluation.

Included in the framework is a GUI that can be run directly on the robot or on a computer connected via WLAN (see Fig. 3). The user interface is realized using Qt4 and OpenGL. This feature has shaped up as a very important tool for understanding and improving the complex algorithms needed for a fully autonomous robot.

The principal screen shows a 3-dimensional view of the internal robot model, the current sensor measurements and the generated map. It allows you to switch into the viewpoint of different sensors and overlay camera images with the 3D view. In addition, various dialogues for starting the different games, monitoring the module activity and CPU usage, calibrating sensors, creating data sets for object and face detection, map editing, navigation and visualizing sensor data can be accessed.

3.3 Simultaneous Localization and Mapping




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

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 a map of its environment in real time during normal operation. The map is based on laser scans, odometry information and (optionally) measurements of the sonar sensors. Fig. 4 show an example of such a map.

The map is stored in two planes. One plane counts how often a cell was surpassed by a laser beam, while the second plane holds information about how often the laser beam was reflected at this position. The occupancy probability for a cell is then calculated as the quotient of the two planes.

3.4 Navigation in Dynamic Environments

In real-life situations, the approach described above 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 enough information.

Thus, our navigation system, which is based on Zelinsky's path transform (see [Zel88,Zel91]), always merges the current laser range scan as unsurpassable object 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 in highly dynamical environments.

3.5 Autonomous Exploration

Several tasks in the @home league, like the “lost and found” game of previous competitions, require the robot to autonomously explore its environment.

For this purpose, we are using a novel exploration algorithm combining Yamauchi’s frontier based exploration [Yam97] with Zelinsky’s path transform. The path transform is extended in a way that not the cost of a path to a certain target cell is calculated, but instead the cost of a path that leads to a close frontier (see fig. 5). More details can be found in [WP07].




Fig. 5. Illustration of the exploration transform algorithm described in [WP07]

3.6 Human-Robot Interface

The robot is equipped with speakers and a microphone, which enables it to communicate via a speech interface. In addition, it has a small screen that will be used to display facial expressions and state information. On the software side, we decided to use two Open Source libraries: *pocketsphinx*³ for speech recognition and *festival*⁴ for speech output.




Fig. 6. Sample output of our hand detection algorithm. Left: 2D visualization with hand regions marked white. Right: 3D visualization of the bounding box of the detected hand volume.

(New in 2010) To enable more sophisticated and robust human-robot interaction, we are currently working on hand-gesture recognition using depth images from our time-of-flight camera. The gesture recognition is a two-step process. First, candidate hand regions are extracted by region growing from possible hand locations, which are expected to be local minima in the depth image.

³ <http://www.speech.cs.cmu.edu/pocketsphinx/>

⁴ <http://www.cstr.ed.ac.uk/projects/festival/>

From those candidates regions, the ones similar to hands are selected. The algorithm is described in [MVGP10] and was submitted for the RoboCup Symposium 2010. In the second step, tracking and interpretation of hand regions is performed. We are currently evaluating an algorithm based on the work presented in [HTH00].

3.7 Object and Face Recognition

The object recognition algorithm we use is based on Speeded Up Robust Features (SURF) [BTVG06], which are local scale-invariant features of grey 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 histogram. This way, sets of consistent matches are obtained. The result is further optimized by calculating a homography between the matched images and discarding outliers.

3.8 Person Tracking

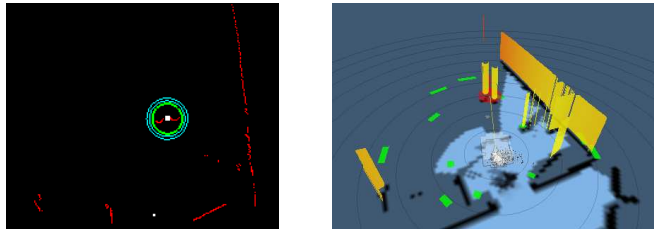


Fig. 7. Left: result of arc segment search in two-dimensional laser range data, Right: GUI visualization during follow mode. The red circle indicates the position of the tracked person.

The people detection and tracking is based on data provided by both LRFs. As the upper laser scans for the upper body of a human, the lower laser is scanning for the legs. Legs are detected by looking for pairs of arc segments in the range data as in [XCRN05]. Then, the presence of the upper body part in the measurements of the upper laser scanner is verified.

3.9 Object manipulation

(New in 2010) For detection and localization of manipulable objects the robot is acquired with a time-of-flight camera and a laser range finder mounted on a pan-tilt unit. While the time-of-flight camera is used for detection of objects on higher locations as tables and shelves, the combination of Hokuyo LRF and pan-tilt unit is favored for objects on the ground. This is due to the laser range finder's higher resolution.

The three-dimensional point cloud which is obtained by either of the above sensors is segmented to isolate objects. Based on their dimensions candidates are selected. A camera is employed to select objects by visual object recognition. Finally, the selected object is manipulated using the 2-DOF gripper or 6-DOF robotic arm.

Traditionally, path planning for robotic arm manipulators is performed in configuration space, because it allows to compute the path between two positions as linear interpolation. However, computation of inverse kinematics is costly and complex. Therefore, we opted for a novel 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 as performing a smooth path or keeping a maximum distance from obstacles using heuristic and cost functions. The path planning approach was developed by [CL10].

4 Conclusion

In this paper, we gave an overview of the approach used by team homer@UniKoblenz for the RoboCup@home competition. We presented a combination of out-of-the box hardware and sensors and a custom-built robot framework. We explained the fundamentals of our message-based robot architecture, the use of well-established techniques like SLAM based on a particle filter and a grid map, and object recognition using matching and clustering of local invariant features.

Based on the existing system from last years competition, effort was put into enhancing the manipulation as well as human-robot interaction abilities of our robot. Three-dimensional scans of the environment are aquired using a combination of a two-dimensional laser range finder and a pan-tilt unit as well as a time-of flight camera. They are employed to detect and localize manipulable objects. A gripper as well as a robotic arm are used for object manipulation. Path planning for the robotic arm is done using a novel graph-based search algorithm based on motion primitives in working space.

References

- [BTVG06] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. *ECCV*, pages 404–417, 2006.
- [CL10] Chitta Sachin Cohen, Cohen and Maxim Likhachev. Search-based planning for manipulation with motion primitives. To appear on ICRA 2010, 2010.
- [HTH00] Pengyu Hong, Matthew Turk, and Thomas S. Huang. Gesture modeling and recognition using finite state machines. In *In Proceedings of the Fourth IEEE International Conference and Gesture Recognition*, pages 410–415, 2000.
- [MVG10] Andreas Muetzel, Sebastian Vetter, David Gossow, and Dietrich Paulus. Hand recognition using time-of-flight camera technology. Submitted for the Robocup Symposium, 2010.
- [WP07] Stephan Wirth and Johannes Pellenz. Exploration transform: A stable exploring algorithm for robots in rescue environments. *Workshop on Safety, Security, and Rescue Robotics*, <http://sied.dis.uniroma1.it/ssrr07/>, pages 1–5, 9 2007.
- [XCRN05] M. Xavier, J.and Pacheco, D. Castro, A. Ruano, and U. Nunes. Fast line, arc/circle and leg detection from laser scan data in a player driver. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation (ICRA)*, 2005.
- [Yam97] B. Yamauchi. A frontier-based approach for autonomous exploration. *1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation*, page 146 ff., 1997.
- [Zel88] Alexander Zelinsky. Robot navigation with learning. *Australian Computer Journal*, 20(2):85–93, 5 1988.
- [Zel91] Alexander Zelinsky. *Environment Exploration and Path Planning Algorithms for a Mobile Robot using Sonar*. PhD thesis, Wollongong University, Australia, 1991.