Interactive Learning of Robot Manufacturing Assistants

Andreas Stopp Sven Horstmann Steen Kristensen Frieder Lohnert

DaimlerChrysler Research and Technology
Cognition and Robotics Group
Alt-Moabit 96A, D-10559 Berlin, Germany
andreas.stopp@daimlerchrysler.com

Abstract

In this paper, research towards interactive learning for robot Manufacturing Assistants is presented. The aim of this research is to develop a robot which can easily be instructed how to either perform tasks autonomously or in co-operation with humans. We describe the prototype of our Manufacturing Assistant and the methods developed for teaching new tasks and environments. The functionality has been demonstrated in a number of factory settings. In this paper, some application examples of our methods are presented.

1 Introduction

In this paper we describe past and ongoing research and development efforts at DaimlerChrysler Research and Technology's Cognition and Robotics Group where over the last years research work has been conducted on human-friendly robots for space, office, and factory automation and towards Manufacturing Assistants.

It is our conviction that the use of mobile robot assistants in manufacturing environments (Manufacturing Assistants) will lead to significant improvements of industrial production processes, particularly in terms of increased productivity and humanisation of the work place. Robot assistants in manufacturing will accomplish tasks through close interaction with people, thus supporting human workers, not replacing them. The human worker is responsible for the command, supervisory, and instructional functions, while the robot assistant will carry out boring, repetitive and strenuous operations. In cases where the robot does not know how to proceed, the human worker will intervene to provide guidance and additional instruction. The robot and the human worker are, therefore, partners in a joint manufacturing process.

Real, complex factory environments are characterised by frequent changes, by varying positions of transport containers, by parts of differing forms and weights in the containers, and by the use of various machining tools. Thus the use of Manufacturing Assistants in real factory environments requires a maximum of flexibility. We believe that this flexibility can only be achieved by instructing the robot assistant in an interactive teaching and learning process.

Therefore, a major goal of our work has been and still is to develop robots that can assist, co-exist with, and be taught by humans. Apart from developing the "standard" mobile robot capabilities such as landmark recognition, path planning, obstacle avoidance etc. our research effort has been aimed at the development of learning capabilities that will allow the user to quickly and intuitively teach the robot new environments, new objects, new skills, and new tasks. In this paper we present some of the results of this work.

Figure 1: Manufacturing Assistant at DaimlerChrysler Research Berlin.

Current research is aimed towards improving the man-machine interaction by adding more advanced communication and cognition capabilities. This has the purpose of further simplifying the teaching of the robot but also to make it more ”co-operative” by having it interpret human commands and behaviour in the given context, allowing it to make better decisions about when, how, and where to assist the human co-worker. An important criterion is, however, that the robot can also perform tasks autonomously once instructed/taught by a human worker. Additionally, it should be able to learn incrementally, i.e. to improve its performance during task execution by ”passively” receiving or actively requesting information from the human (the latter could for example be in the case where the robot detects ambiguities which it cannot autonomously resolve).

Thus, in some sense we want to create a life-like assistant, but the fact that this assistant should be able to communicate and cooperate with the human worker, i.e. not merely be an idiot savant, means that the internal representations used by this assistant should be symbolic to ensure that these can be inspected, interpreted and possibly adjusted by the human worker. This may be a significant difference to traditional Artificial Life research where focus tends to be more on learning “biological”, sub-symbolic representations.

2 Manufacturing Assistant

The DaimlerChrysler Manufacturing Assistant is a modular arm/platform system fit for industrial use. A first prototype is shown in Figure 1. Its features are: a multi-skill oriented system and control architecture, multiple sensors for interaction; 2D/3D laser range sensors, vision (gripper camera), force torque sensor, 6-DOF-Mouse, pentop PC and headset as multi-modal commanding device (MMI, see Figure 2).

A typical scenario for a robotic assistant in an industrial setting would be:

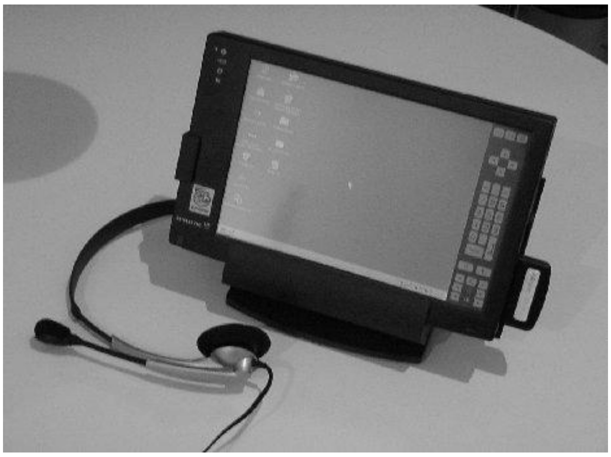

Figure 2: Pentop with headset serving as portable MMI for the Manufacturing Assistant.

- the robot is led through the factory halls and is shown important places (stores, work stations, work cells etc.),
- the robot is shown relevant objects, e.g. tools, work pieces, and containers,
- the robot is shown how to dock by work cells, containers etc. in order to perform the relevant manipulation tasks,
- the robot is taught how to grasp various objects and how (and possibly in what sequence) to place them in corresponding containers or work cells
- in case of a co-operation task, the robot is shown when and how to assist the human worker.

The Manufacturing Assistant approach is related to the Cobot approach [1, 3] but where Cobots are active-passive devices, implementing virtual walls, virtual fixtures etc. to guide and help the human worker, Manufacturing Assistants are equipped with active degrees of freedom and have a higher degree of autonomy.

3 Architecture

Our scalable architecture for future robotic assistants is a multi-skill oriented system and control architecture. Thus, the robot can extend and adapt its skill profile most appropriate to the current task either task/plan-driven or event-driven. Figure 3 shows the generic scheme of our architecture for future Manufacturing Assistants. The lower part shows the components of the complex control loop of a modular arm/platform system and the upper part shows the interfaces and the higher-level methods. The robot control loop is divided into the control loops for the manipulator(s) and for the platform. Both have their own sensors, recognisers, low-level controllers and actuators. Perception and coordinated robot behaviour are organised by task-oriented sets of skills supported by access to extensive world model data (knowledge) and by higher level planning and learning methods.

For human-robot interaction a multi-modal commanding device (pentop-PC and headset) is connected to the Manufacturing Assistant via radio Ethernet. Additionally, for specific applications, pointers, laser pointers and gestures can be used for commanding and teaching in connection with

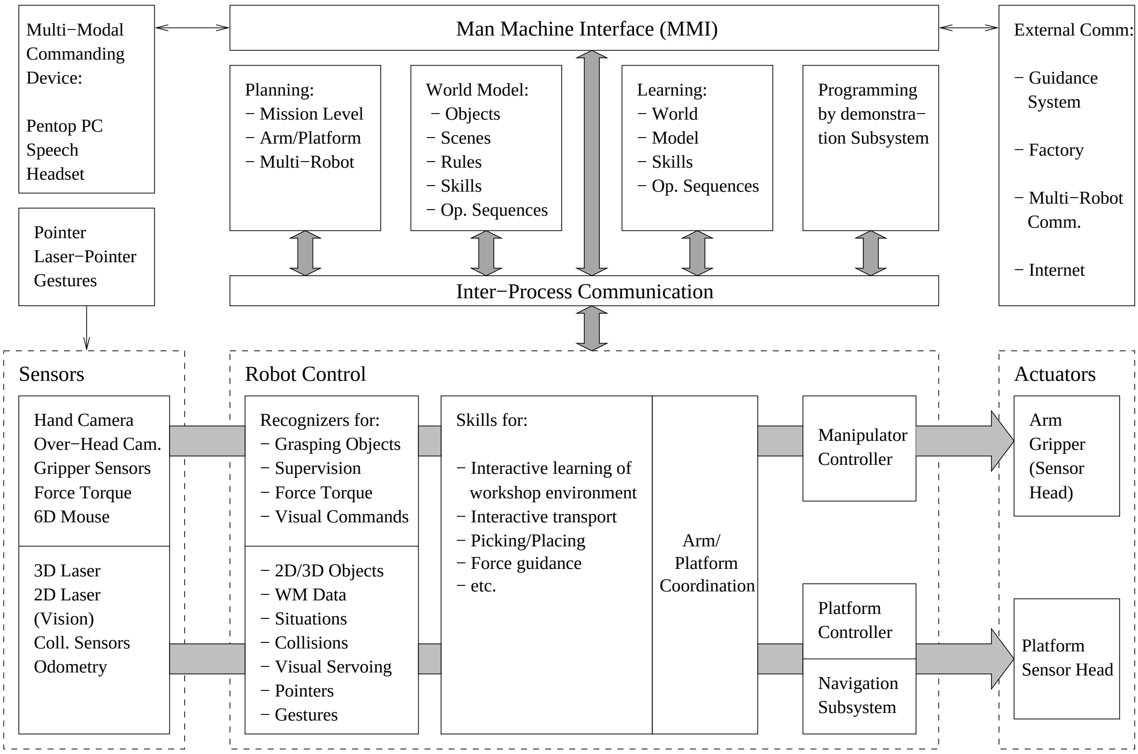

Figure 3: Manufacturing Assistant: Software architecture of a modular arm/platform system.

available sensors and appropriate recognisers. For providing our Manufacturing Assistants with a maximum of knowledge, interfaces to the workshop environment, to cooperating robots and to the Internet (e.g. for diagnostics, maintenance, remote control or external knowledge access) are installed. Our software is running on multiple on-board industrial PCs using the real-time operating system QNX.

4 Technologies

As described above, future robotic assistants require innovative methods for real intelligent and cooperative robot behaviour. The research work of our Lab is focused on 2D and 3D recognition and scene analysis using laser sensors, situation analysis for safe interaction, hybrid reactive planning methods for motion and manipulation, fast multi-sensor-based control, and new learning principles for interactive teach-in and sensor-based learning. Dependability issues are of course also of paramount importance to this research field [7] but will not be treated further in this paper.

5 Interaction Methods

As stated in the introduction, we believe that the only feasible way to endow the Manufacturing Assistant with sufficient flexibility to deal with the very diverse and dynamic tasks in a production context is to make it capable of easily learning new environments, objects, skills, tasks etc. In the following sections we provide a few examples of how the robot is taught interactively in its "natural" environment.

Figure 4: Initial learning of an workshop model by our autonomous robot platform—starting the learning procedure by significant gestures and following the human for learning.

5.1 Learning the Environment

Although scientifically challenging, autonomous exhaustive exploration is not appropriate in real manufacturing environments. This has several reasons:

- Manufacturing environments are often quite open, i.e. there are not always physical barriers separating the areas where heavy machinery like a fork lifter is allowed to go and where not.
- It is in general not safe to have an industrial strength mobile platform without a map of the environment moving un-supervised through the environment.
- Due to the dynamics of the environment (areas temporally blocked by containers, other platforms etc.), it is not possible within a realistic time period to guarantee that the robot has autonomously explored all the places relevant for its operation.
- In order to be able to communicate with the human workers the Manufacturing Assistant must know what the various areas and work stations are called which is only possible if some operator teaches this information. The most intuitive and robust way to do this is on-line and on-site.

We have therefore chosen to teach the initial environment model using human guidance for focusing attention to accelerate the learning procedure. The idea is to lead the Manufacturing Assistant around in the relevant part of the factory using a few simple but robust gestures (see Figure 4). We would like to stress that the goal is not to explicitly teach the robot all features of the environment but to show where it should itself generate its environment model. When the robot is performing its tasks, this model, containing only the major structures of the environment, is on-line extended with current local information enabling the robot to adapt to the actual situation.

The model is generated by scanning the part of the environment shown by the human operator using a 2D laser-scanner. Through the movement of the robot, a 3D point image is created which

Figure 5: The worker is pointing out objects directly in the scene using a laser pointer.

is subsequently segmented into planar patches which are projected onto the ground plane. This way a 2D map of the environment is created which is efficient for the robot planning process in terms of memory and computation time consumption. In the map, information about unseen areas is explicitly stored which means that the robot, on subsequent visits to the area, can autonomously extend this in case it should be able to scan previously unseen areas. Due to the intuitive, symbolic representation it is furthermore easy for the human operator to verify the map and for example to extend it with virtual walls.

5.2 Learning Objects

The general idea we pursue for interactively teaching objects is to let the operator point out the relevant objects/features either directly in the world (e.g. by using a laser-pointer, see Figure 5) or in a graphical interface showing the relevant (possibly pre-processed) sensor data. In Figure 6 is shown an example of teaching objects to be grasped from a conveyer belt using a gripper camera. With this system, developed by Graphikon GmbH as a part of the MORPHA Project, the user simply places the relevant objects under the camera a few times, showing various examples (of the same aspect). For each example, the object is pointed out in the image (in the teaching phase, objects have to be non-overlapping) and at the end, when sufficient examples have been taught, a grasping position is defined. This is quite intuitive and has been proven to work reliably under real world conditions. The representation actually learned by the system are means and variances of various object geometry parameters which are used as bounds for a standard box classifier. This has the advantage that these can be interpreted by the user for whom is it then possible to query about the reasons in case for instance a specific part has not been correctly recognised.

We have developed similar methods for teaching 3D objects in laser data [4]. An example of this is the teaching of polyhedral objects which is accomplished by presenting a plane-segmented 3D image to the user, who can then simply point out the surfaces belonging to the object he/she wants to teach. Similarly, free-form objects can be taught by simply cutting the relevant part out of a 3D point cloud acquired by the 3D laser-scanner. These models are e.g. appropriate for use with the iterative closest point (ICP) algorithm [2].

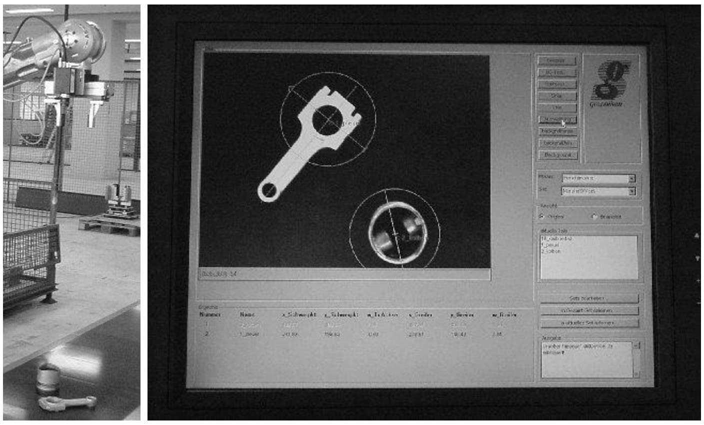

Figure 6: Manufacturing Assistant: learning objects. Left the scene with the manipulator looking down on the conveyer belt with the objects. Right the GUI used to teach the objects.

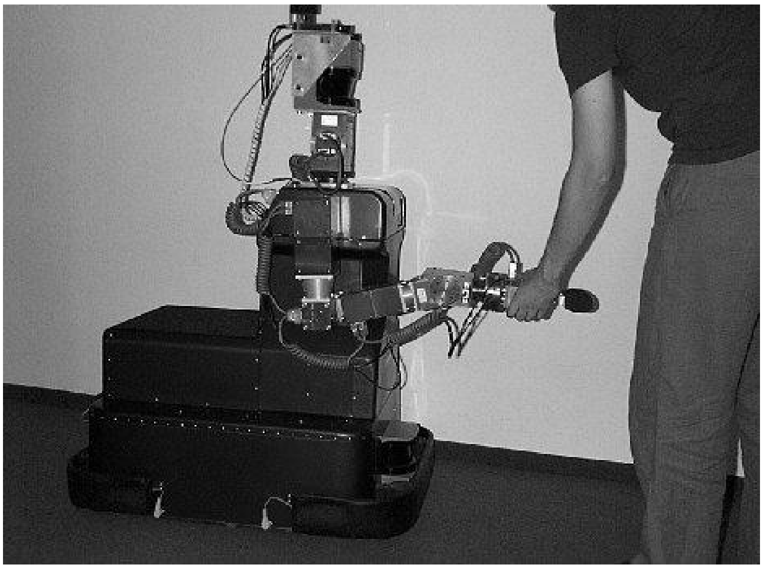

Figure 7: The Service Robot "Clever" at DaimlerChrysler Research and Technology, Berlin.

5.3 Interaction Methods for Learning Skills

While traditional robot systems must be programmed using teachboxes or control-panels, we chose to use more intuitive ways for teaching arm and platform movements, which can not only be reproduced, but also be adapted to slight changes in the environment or situation. We use the force-torque sensor on the wrist to detect forces applied by the user, resulting in a coordinated motion of arm and mobile platform, in which the arm's motion compensates for the holonomic constraints of the platform [6]. We use this mode of interaction to teach pick-up- and place-movements and platform paths (see Figure 7). We even employ this strategy for hand-eye calibration, by leading the gripper to a previously visually recognised target. The coordinated arm/platform-movement has been implemented on our service robot demonstrator "Clever" [5], but will be transferred to the industrial manipulator as soon as its firmware allows us to establish the required control loops.

In general, the arm movements are taught in a sensor context where the object in question initially has its type and position relative to the robot determined from sensor data. This means that the arm trajectories can be explicitly stored (and thus visualised) in a coordinate frame relative to this specific type of object, which allows for easy generalisation to situations where the object is recognised at some other position. In a cooperation context, this is a clear advantage over typical visual servoing techniques, where a control law is defined or learned which typically uses a few low–level features to guide the robot gripper to the desired position. In return, the latter approach has the advantage of implementing a control loop which is tighter than is typically possible with our approach which in principle requires the object to be recognised in each control cycle. For applications where the objects are static or where their movements are known (e.g. as when they move on a conveyer) this is not a practical problem, though.

6 Conclusion and Further Work

The successful use of robot assistants in a manufacturing environment will depend on the critical question as to whether the share of the work done by the robot assistant is as high as possible, and whether it is capable of flexibly *and safely* dealing with varying sequences of tasks under variable boundary conditions.

To achieve this, we think it is crucial that the robot is capable of one-shot learning and to autonomously adapt to the actual conditions. To maximise safety and robustness, we want the on-line adaptations to be as small as possible, which poses further demands to the learning process. We believe that the only way to achieve this kind of high-quality, one-shot learning is to use human guidance. In this paper we have presented various techniques we have investigated for teaching different kinds of model information to a robot working in an industrial environment.

Although interactive teaching and learning by demonstration are not typical themes for Artificial Life, we think these techniques deserve some attention for transferring skills between higher level agents of different kinds, which may often not be able to directly exchange internal representations.

The limitations of using exchangeable, high-level type representations is that the various agents must possess an architecture which knows how to deal with these representations and how to derive from them the actions eventually leading to the desired outcome. This clearly poses an overhead in comparison to systems where e.g. control laws (perception-action couplings) or classification information is learned directly on a more “embedded”, sub-symbolic level.

We are currently integrating the presented techniques on the Manufacturing Assistant.

7 Acknowledgements

This research was partly sponsored by the German Ministry for Education and Research under the projects NEUROS, Neural Skills for Intelligent Robot Systems, and MORPHA, Intelligent Anthropomorphic Assistance Systems.

References

[1] P. Akella, M. Peshkin, E. Colgate, W. Wannasuphoprasit, N. Nagesh, J. Wells, S. Holland, T. Pearson, and B. Peacock. Cobots for the automotive assembly line. In *Proceedings of the 1999 IEEE International Conference on Robotics and Automation*, pages 728–733, Detroit, Michigan, May 1999. IEEE.

[2] P.J. Besl and N.D. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.

[3] R. B. Gillespie, J. E. Colgate, and M. A. Peshkin. A general framework for cobot control. *IEEE Trans. on Robotics and Automation*, 17(4):391–401, 2001.

[4] Steen Kristensen, Volker Hansen, Sven Horstmann, Jesko Klandt, Konstantin Kondak, Frieder Lohnert, and Andreas Stopp. Interactive learning of world model information for a service robot. In H. I. Christensen, H. Bunke, and H. Noltemeier, editors, *Sensor Based Intelligent Robots*, Lecture Notes in Artificial Intelligence (1724), pages 49–67, Berlin, 1999. Springer.

[5] Steen Kristensen, Sven Horstmann, Jesko Klandt, Frieder Lohnert, and Andreas Stopp. Human-friendly interaction for learning and cooperation. In *Proceedings of the 2001 IEEE International Conference on Robotics and Automation*, pages 2590–2595, Seoul, Korea, 2001. IEEE.

[6] Steen Kristensen, Mathias Neumann, Sven Horstmann, Frieder Lohnert, and Andreas Stopp. Tactile man-robot interaction for an industrial service robot. In H. I. Christensen and G. Hager, editors, *Sensor Based Intelligent Robots*, Lecture Notes in Artificial Intelligence, Berlin, 2001. Springer.

[7] Y. Yamada, T. Morizono, and Y. Umetani. A consideration toward human/robot dependability based on the current techniques of securing human safety for human/robot collaborative conveyance tasks. In *Proceedings of the 2001 IARP/IEEE-RAS Joint Workshop on Technical Challenges for Dependable Robots in Human Environments*, Seoul, Korea, May 2001.