

The Manufacturing Assistant: Safe, Interactive Teaching of Operation Sequences *

A. Stopp T. Baldauf R. Hantsche S. Horstmann S. Kristensen
F. Lohnert C. Priem B. Rüscher

DaimlerChrysler Research and Technology, Cognition and Robotics Group
Alt-Moabit 96A, D-10559 Berlin, Germany

E-mail: {thomas.baldauf, sven.horstmann, steen.kristensen, frieder.lohnert, christopher.priem, bjoern.ruescher, andreas.stopp}@daimlerchrysler.com

Abstract

In this paper we describe our Manufacturing Assistant, a mobile manipulator for future flexible manufacturing, and its capabilities for interactively learning new operation sequences. As an example we present a method for interactively teaching an order-picking task. The task is taught using a laser pointer and a hand-held computer. Furthermore, the system interacts with the user using speech output. The system was successfully demonstrated on the Hannover Industrial Fair 2002.

1 Introduction

The aim of our work on the Manufacturing Assistant is to develop a system, with which a human worker and an industrial robot can work as partners in a joint manufacturing process and so combine the intelligence and cognitive capabilities of the human with the strength, speed and precision of the robot to optimise production in terms of quality and flexibility [5, 6]. In current manufacturing systems, humans and robots are strictly separated due to safety issues and because the robots lack the capability to react to and learn from the humans, hence making human presence counter-productive. Thus, the challenge of our task was twofold: to create a system where the human cannot be subject to danger due to the robot, and where the robot is capable of interacting with the human in order to learn new or to improve existing tasks.

As an example of an operation sequence demonstrating the cooperation of human and robot we chose order-picking, i.e. the task of picking and placing a

number of possibly different objects on a common carrier, since this is a realistic, labour-intensive task often encountered in manufacturing and in logistics.

To comply with the existing safety regulations for industrial manipulators we decided to use certified safety sensors to monitor the robot work space and to let the human operator command the robot from outside its work space using a laser pointer and a hand-held computer connected to the robot over a wireless LAN.

The learning part is divided up into more parts; learning the relevant objects, learning how to grasp them, learning how many of each type to use, and finally learning where to place them on the carrier. Each of these steps is described in detail in the following sections.

2 The Manufacturing Assistant

The Manufacturing Assistant (see in Figure 1), on which the methods were developed, consists of an arm module with a 6DOF Reis industrial manipulator with a gripper camera and on-board computers for image processing and dynamic path planning and control. The arm module can be moved around by the autonomous forklift truck. This feature was, due to space limitations on the Hannover Fair, not used in the order-picking task described in this paper, though. An important difference to the Cobot approach [1, 4] is that 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.

It is our conviction that the use of mobile robot assistants in manufacturing environments (Manufacturing Assistants) will lead to significant improvements

*This research was partly sponsored by the German Ministry for Education and Research under the project MORPHA, Intelligent Anthropomorphic Assistance Systems.




Figure 1: *The DaimlerChrysler Manufacturing Assistant.*

of industrial manufacturing processes, particularly in terms of increased productivity, flexibility 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.

In a previous paper [8] we have shown how the Manufacturing Assistant can quickly learn new environments, objects and skills using human-robot interaction.

In this paper we want to show how it is possible to safely and interactively teach the Manufacturing Assistant an order-picking sequence which it can, commanded by the factory production system, autonomously repeat as many times as needed. Previous to teaching the order-picking sequence itself, the system must know where the containers with the various objects are, and what these objects look like.

This information is assumed available from the factory production system or from previous teaching procedures [8]. The action sequence for teaching the order-picking task is:

1. The operator specifies via the hand-held computer in what sequence how many objects of each type should be placed on the carrier.
2. For each object, the robot positions the gripper camera over the carrier and the user is prompted to indicate with the laser pointer where he/she wants the next object to be placed.
3. The user points to the desired position with the laser pointer and triggers the laser point acquisition with a wireless trigger integrated in the pointer.
4. The system checks that the object can be placed at the shown position without causing collisions with objects already present on the carrier. If a possible collision is detected, the user is informed about the problem and is asked to enter a new position. Otherwise, the robot moves to the container with the relevant object type, recognises one, picks it up, and places it at the desired position.



Figure 2: The setup for the order-picking demonstrator shown at the Hannover Fair 2002. The operator is showing with the laser pointer where the robot should place the next item. The items are retrieved from the storage areas left and right to the robot. The striped line in the floor marks the border which, when unexpectedly crossed, will cause the robot to come to an emergency stop. This is ensured by the Sick PLS laser scanner partly visible in the upper right corner of the image.

5. The system records the sequence and the taught object positions on the carrier and can thus repeat the task autonomously.

The setup for the order-picking task and the corresponding functional block diagram can be seen in Figures 2 and 3.

3 Object and Laser Point Recognition

The object recognition module, which can also recognise laser points, was provided by the computer vision company Graphikon GmbH who is a partner in the MORPHA Project. The system is commercially available and will thus not be discussed in detail here. A special feature about this module is its learning capability which enables it to learn new objects simply by placing them under the camera in a few different poses, see Figure 4. This makes it quite easy to add new object types to the system. Furthermore, the appropriate grasping point and orientation can be shown, which are then output along with the object type when an object is recognised.

The hand-eye calibration, i.e. the process of establishing the gripper to robot relation, is performed interactively by first letting the object recognition system recognise a set of calibration points on a calibration object. Afterwards, the robot's TCP is moved

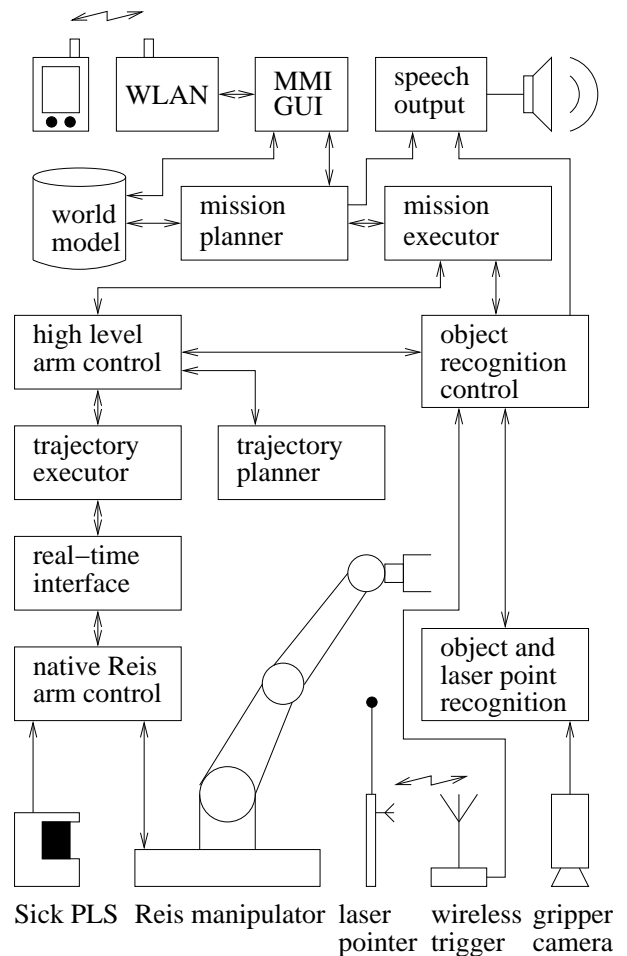


Figure 3: Functional block diagram of the Manufacturing Assistant's arm module.

to each of these points using the Reis manipulator's capability to be directly touched and guided using a 6DOF mouse. Thus having a number of sets of corresponding image and robot TCP coordinates, the transformation can be derived. The calibration need only be performed whenever the camera's position on the robot has changed.

4 Mission Planning and Execution

The mission planning and execution modules are central to the system since they receive the commands from the user (or the factory production system), plan the corresponding actions to be carried out by the system, and subsequently monitor that everything is executed according to that plan.

The hierarchical planning and execution system (HIPLANEX) is a STRIPS-like forward symbolic planner [2], which can handle the extended ADL (Ad-



Figure 4: 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.

vanced Description Language) features like universal quantification and conditional effects. It uses a best-first strategy searching the state space for a sequence of missions to achieve the given goals. Furthermore, a conditional planner [3], which may generate different plans for different possible values of plan-states, is integrated.

The idea of hierarchical planning [7], where the overall plan is generated by successively refining initial, abstract plans, is to reduce the complexity of the plan generation process and allow a more human-like approach to mission planning. The process of reducing abstraction ends with only having elemental, directly executable missions in the plan. The abstraction levels of available missions are defined manually inside the mission definitions.

The planning challenge of the order-picking task stems from the required interlocking of planning and execution. This interlocking is due to the fact that the object positions (in the container as well as on the carrier) and in some cases the number of objects are not known at plan time since these are provided by the operator and by the object recognition system. This is dealt with by using run time controlled loops, by employing generic objects with pseudo parameters at plan time which are replaced at run time, by starting different plans within loops of execution, and by using an inventory management for the objects and their positions.

Two high-level missions “OrderPicking” and “CarrierClear” were defined for placing any number of objects from their (different) containers to one or more carriers and placing them back to their original positions (the latter mission was defined to make it possible to repeat the demonstration arbitrarily often). The operator selects the missions and their parame-

ters via a man-machine-interface, in case of “OrderPicking” the number of objects and their destination carrier(s) and in case of “CarrierClear” only the carrier(s), which should be cleared. When the object selected by the operator is a concrete instantiation of an object, known in the world model—as is the case with the carriers—this instantiation is directly used for the planning. Often, however, the operator rather selects a type of object which is to be placed on the carrier. In this case, a generic object “prototype” is used which e.g. has no specific position but only a reference to the (known) container, in which it can be found. When such a generic object is selected, the planner automatically incorporates an elemental mission for recognising its position at execution time.

Since the teaching phase is executed only once and afterwards the mission should be repeated autonomously several times with the same deposit-positions (the number of repeats can be specified within the MMI), the laser point recognising missions are needed only during the first execution and have to be skipped otherwise, which is handled in the executor.

To avoid using object and position parameters in the “CarrierClear” mission, these are stored internally in the database during the “OrderPicking” task using an rudimentary inventory management. The stock data is also accessible from inside the mission definition so it can be used for planning purposes. To allow the planning of “CarrierClear” without first having executed “OrderPicking” (e.g. to first fill up the stocks), the number of objects to be (re-)placed is not retrieved from stored stock data. Instead a conditioned execution loop is executed, which terminates when no further object is found (recognised) within the carrier. This conditioned loop is directly specified in the mission definition.

The plan and execution results are displayed in the GUI, where the known objects, containers and positions are represented as polygon shapes in a 2D-coordinate system. The creation of objects out of prototypes and their positional changes after execution of each plan step are displayed. This allows for monitoring the planning as well as the execution phase.

5 Object Recognition Control

The object recognition control module receives symbolic commands like “recognise deposit pose for piston on carrier2” or “recognise piston in box4” from the HIPLANEX system. In case a generic object like a piston must be recognised the controller activates the object recognition system and waits for the results. If the object recognition system does not recognise any

objects from the gripper camera's current viewpoint, the controller calculates a scan path from knowledge about the camera position and the container object (e.g. "box4") and issues corresponding commands for the robot control system to move the gripper camera. If the scan is completed without any objects of the correct type being recognised, an error message is returned to the mission executor and a corresponding text message is sent to the speech output module that thus informs the operator about the error.

When the object recognition control module is told to recognise laser points it likewise activates the object recognition system which will (if the operator is pointing to the carrier) send a continuous stream of recognised points to the controller. The controller, on the other hand, waits for a trigger signal from the operator, who thereby indicates that the laser point is now at the desired position. This causes the controller to sample the current position and return it to the mission executor. However, prior to this, the object recognition controller checks if the object can be placed at shown destination pose without it or the gripper colliding with already present objects. For this operation, the controller uses the information from the inventory management and knowledge about the object to be placed on the carrier. If the pose is not valid, the operator is informed about the problem and asked to select a new position via the speech output module.

6 Manipulator Control

The manipulator control part of the system comprises the high-level controller, the trajectory planner, the trajectory executor, and a real-time (interpolation-rate) interface to the native robot arm control unit allowing the implementation of e.g. sensor-based feedback loops. This interface is a commercial product from Reis Robotics and will thus not be described further here.

The high-level controller coordinates the planning and execution of symbolic movement commands coming from the mission executor and from the object recognition control.

The trajectory planner was created especially to account for the fact that it must be able to quickly generate new trajectories and at the same time to provide a maximum of safety and "confidence" for the human operator. We therefore decided to use a graph-based planner which, as far as possible, plans trajectories on a 3D-graph surrounding the robot. The idea is that the graph establishes a "highway" net on which most of the TCP "driving" is done. So when the arm must move from its current position to some point,

A , it searches the nearest node in the graph (in the direction of A), plans a path to that node, moves on the graph to the node nearest to A , and plans the remaining part from here to point A . Apart from being quite fast, this has the advantage over e.g. stochastic planners that the planner has a deterministic output, which is important for certification, and that the arm largely moves in the same ways which we have found makes it easy and fast to see, if the manipulator is doing what was intended and if it is operating correctly.

The graph is currently generated manually, but we expect that it will be possible to automatically adjust the graph, e.g. according to where the containers with parts are located.

7 Safe Physical Interaction

Under some circumstances it may be desired that the robot and the human operator physically share the work space. This may be the case if the operator wants to correct the robot or if he/she must perform some operation inside the robot's work space. In the latter case, this may be a fixed part of the plan, i.e. that the operator due to his/her superior cognition and dexterity performs part of a joint human-robot assembly task.

On the Hannover Fair we demonstrated this kind of planned human interaction in the robot's workspace by letting the operator emulate the conveyer belt present in our normal setup (see Figure 1) but due to space limitations not on the Hannover Fair. Thus, to be able to show the autonomous repetitions of the taught order-picking task, a human operation was included in the plan. Of course by such operations the major concern is human safety and it was therefore chosen to rely on certified safety mechanisms already present in the employed systems.

In the normal operation mode, a Sick PLS laser scanner is surveying the open side of the robot's "cage" to ensure that nothing crosses this border. Any violation produces an emergency stop of the robot. When the operator is supposed to enter the work space, however, the robot is first driven to a fixed (but programmable) "safe" position where the Reis robot's safety controller is guaranteed (and certified) to perform an emergency stop in case the robot leaves it. Then the Sick "laser curtain" is opened and the operator is via a light signal and speech output requested to perform the relevant operations. When the operator is done, he/she confirms this by activating a key switch upon which the Sick "curtain" is again closed and the robot continues its work.

It is also possible for the human to directly move the robot arm using a 6DOF mouse, but due to the

quite slow TCP speeds at which this is allowed, we find this to be most useful for teaching and calibration and in the situations where the manipulator is used as a power amplifier to handle heavy parts.

However, for interactively performing planned, routine operations we prefer to either teach or command the robot from outside its working space using the laser pointer. Note that although commanded by the operator, it is still possible for the robot to supplement the given information using its own sensors e.g. in order to exactly determine where to grasp a given object, how to avoid an obstacle etc.

8 Demonstration Experiences

As previously mentioned, the order-picking system was demonstrated on the 2002 Hannover Industrial Fair in Germany. The fair was running over 6 days and an estimated 100 demonstrations were given. The system ran stably and never failed to recognise laser points and objects nor did it fail to correctly grasp and deposit those.

The fair visitors, who can generally be described as industry automation professionals (but not necessarily all of them robotic experts), found it easy to use the system for teaching the order-picking tasks. Using the laser pointer with the wireless trigger thus proved to be an intuitive and robust way to do human-robot interaction.

9 Conclusion

The aim of this work was to show how humans and industrial robots can cooperate in a joint work process in order to optimally benefit from the strengths of the respective agents. The demonstration scenario was order-picking, where the human with a laser pointer and a PDA teaches the robot the sequence in which a number of objects must be placed on a carrier. Also it was demonstrated how the human can safely perform planned operations in the robot's work space.

Important to us was to show that such a cooperation task can be performed safely and in a not too distant future be introduced into real factories. It was therefore chosen to use standard hardware components and as far as possible robust, well-known data-processing algorithms.

We think that the successful demonstration proves that it is indeed possible, using the right choice of hardware and algorithms, to have humans teach and co-exist with robots in a joint manufacturing process. It is however also clear that increasingly complex processes will demand more work on robot perception and planning if the robot should be a true assistant to the human. Future work will therefore be aimed towards

extending the current sensing and planning capabilities of the Manufacturing Assistant and to improve the work space surveillance to allow for more flexible but safe man-robot interactions.

References

- [1] P. Akella, M. Peshkin, E. Colgate, W. Wannasuphprasit, 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] Eyal Amir. Planning, AI Qual Summary. 1997.
- [3] D. Draper, S. Hanks, and D. Weld. Probabilistic planning with information gathering and contingent execution. In *Proceedings of the 2nd Intl. Conf. on Artificial Intelligence Planning Systems 1994 (AIPS-94)*, 1994.
- [4] 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.
- [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 G. D. Hager, H. I. Christensen, H. Bunke, and R. Klein, editors, *Sensor Based Intelligent Robots*, Lecture Notes in Computer Science, LNCS 2238, Berlin, 2002. Springer.
- [7] I. R. Nourbakhsh. Using abstraction to interleave planning and execution. In *Proc. of the Third Biannual World Automation Congress 1998, (WAC98)*, Anchorage, 1998. TSI Press.
- [8] Andreas Stopp, Sven Horstmann, Steen Kristensen, and Frieder Lohnert. Towards interactive learning for manufacturing assistants. In *Proceedings of the 10th IEEE International Workshop on Robot-Human Interactive Communication (ROMAN'01)*, pages 338–342, September 2001.