

Three Types of Learning Robots

Volker Graefe and Rainer Bischoff

Bundeswehr University Munich, Faculty of Aerospace Engineering
Intelligent Robots Lab (formerly: Institute of Measurement Science)
85577 Neubiberg, Germany, www.unibw-muenchen.de/campus/LRT6

Abstract

Compared to even the most advanced robots existing today, future service and personal robots will need much more intelligence, robustness and user-friendliness. The ability to learn contributes to these characteristics and is, therefore, becoming more and more important. Numerous varieties of learning by robots exist. We have implemented some of them on three generations of autonomous robots and discuss them here together with results of real-world experiments: (1) the acquisition of map knowledge by a vision-guided autonomous indoor vehicle and by a humanoid service robot, allowing them to navigate in typical office buildings, (2) the acquisition of motion control knowledge by a vision-guided calibration-free manipulator, allowing it to gain task-related experience and to improve its manipulation skills while it is working, and (3) the ability to learn how to perform service tasks to the satisfaction of initially unknown humans through situated spoken dialogues with them, thus displaying intelligent and cooperative behavior.

1 Introduction

Intelligence and Learning

Present industrial robots usually must be carefully instructed, programmed and calibrated by specialized experts when they are first installed, and also later, whenever the task or the environment changes. Such an approach is quite unacceptable for the versatile service and personal robots that we hope to have in the future and which we expect to relieve us from all kinds of daily chores in our homes and offices. For them the individual services of expert engineers will be unaffordable and they will have to co-exist and cooperate with humans who are not knowledgeable about, and not interested in, robotics. Similar to humans and animals, those robots will have to acquire by learning the knowledge they need for surviving, performing their tasks, and becoming valuable servants and companions of their owners.

Learning is closely related to intelligence. There are many kinds of intelligence, including the intelligence that enables human experts, for example, to prove theorems or to play chess. In robotics we are mainly interested in a totally different kind of intelligence, a more practical one, like the one that enables, for instance, animals and small children to orient themselves in their environment and to move in a purposeful and goal-directed fashion (cf. [Beer 1990]). Included in the practical robot intelligence

that we have in mind is the ability to acquire necessary knowledge and abilities by learning, i.e. by interaction with the environment. Just as there are many kinds of intelligence, there are also many varieties of learning. Some of them that we have found useful for autonomous robots are introduced in the sequel.

Inspiration by Nature

Living organisms are the most robust and sophisticated autonomous systems that exist. Therefore, in robotics we should sometimes look at nature for inspiration. When an animal is born, it has inherited a repertoire of behaviors and abilities essential for its survival. As it grows older, it improves and refines these behaviors and abilities, and acquires additional ones, by copying, and by trial and error. It also acquires the knowledge necessary for its survival, e.g., about the topography of its environment.

We propose a similar approach in robotics: a robot should be able from the beginning to execute certain behaviors; by learning, it should then improve and refine its abilities and develop additional ones.

Three Types of Robots

As intermediate steps on the way towards realizing truly intelligent robots we have during the past 15 years worked with three autonomous robots and performed real-world experiments with them. They all use vision as their main sensor modality, and each one of them implements at least one specific form of learning. The first one is a mobile robot able to navigate in networks of passageways in ordinary office buildings. The second one is a stationary manipulator that needs no quantitative model of itself, but learns automatically in the course of its work how to control its motions more and more skillfully and efficiently. The third one is a humanoid service robot combining the abilities of locomotion, navigation and manipulation. Moreover, it is able to conduct dialogues with humans in

natural language and to learn through such dialogues a user's wishes as well as facts relating to the user and to the environment, even if the user expresses them in ambiguous ways, as humans often do when they speak naturally.

2 Learning to Navigate

Attributed Topological Map

The basic principle of behavior-based navigation (and of behavior-based robot control in general, as it is now widely used) is the achievement of a desired task by acti-




Figure 1: Section of an attributed topological map. Attributes (here shown in an abbreviated form) with hints on geometrical quantities, visible landmarks (lm), etc. may be attached to each directed path and each named point.

vating an appropriate sequence of elementary behavior patterns. Examples of such behavior patterns in the case of navigation in buildings are “following a hallway”, “turning at an intersection”, or “moving towards a landmark” [Wershofen, Graefe 1992].

In order to orient itself in a network of corridors the robot must have some knowledge of the topology of the passageway network

and it must know how to recognize specific locations in the environment. This knowledge is stored in a topological map which is extended by suitable attribute lists to facilitate reorientation of the robot and coordination of behavior patterns (Figure 1). Intersections and junctions (A, B, C ...) of passageways, and task relevant locations, such as docking stations (e.g., G), are represented as points. As attributes they carry lists of all directly adjoining intersections and junctions, and of the approximate angles between the passageways leading into the point, and in addition, a list of presumably visible landmarks with clues as to how they can be recognized. Depending on the direction in which a passageway is being followed, different landmarks will be visible and, therefore, each physical passageway is typically connected to two attribute lists. If such an attributed topological map is to be used for behavior-based navigation, large errors regarding the exact metric of the modeled network are permitted. It suffices then to represent the topology exactly, since navigation could rely on visual cues only. Such maps are very similar to the subway maps that are available in most large cities. Both kinds of maps are geometrically inaccurate, but since they reflect the topology of a network correctly and indicate the relevant attributes of each line and each station, they are an excellent basis for planning and executing trips.

The characteristics of the attributed topological map required by the behavior-based approach make it much easier to generate such a map than the geometrically accurate map required by some other navigation methods (e.g., coordinate-based navigation, cf. [Graefe, Wershofen 1991]).

Map Building

Two forms of exploring an environment may be considered for generating an attributed topological map automatically by a learning robot, supervised and unsupervised learning. Having the robot explore the environment under the guidance of a human has advantages. The human guide may restrict the exploration to those parts of the environment that are of actual interest, and he may inform the robot of the names that people use in referring to specific locations and landmarks. If needed, a more comprehensive map may be generated later by continuing the exploration.

The operator guides the robot through the network of passageways by deciding at each intersection or other point of interest on the appropriate behavior pattern, such as following the hallway, memorizing the appearance and relative location of a landmark, or turning by a certain angle.

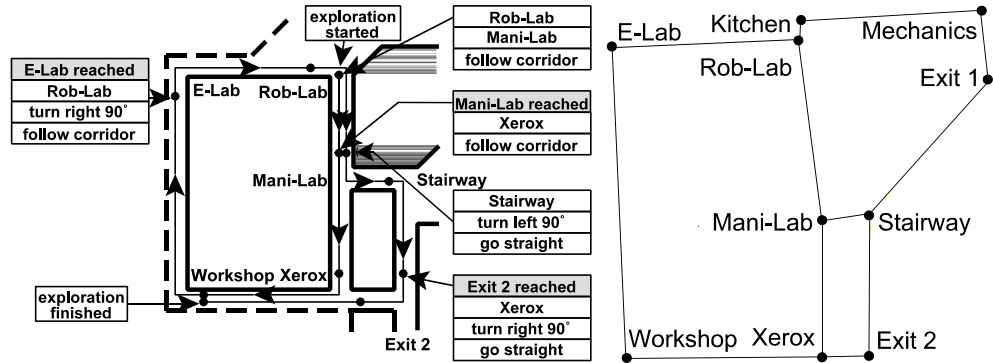


Figure 2: Left: Guided exploration of a network of passageways, Right: Graphical representation of the attributed topological map generated in this exploration and an additional later exploration (upper right part)

During this guided exploration the robot uses its odometry-based dead-reckoning navigator to automatically determine the geometrical attributes (distances and angles). Other map attributes are determined interactively with the operator. At each relevant point the robot asks the operator for the name of the point and enters it into the map. Also, the robot may ask the operator to point out objects that may be used as a reference for executing the indicated behavior pattern. If the robot is able to recognize the proposed object, it generates an internal description and enters it into the map. If not, it requests the operator to suggest a different object.

Adaptation and Incremental Learning

If the robot traverses the same path repeatedly, it may update the attributes that are recorded in the map. This allows errors that may have been made during the initial learning process to be corrected, and it lets the robot adapt to changes in the environment. Moreover, the structure of the map makes it easy to expand the area covered by the map step by step without erasing or forgetting previously learned facts.

Real-World Experiments

The acquisition of map knowledge by supervised learning was developed and evaluated with our mobile robots *ATHENE* [Kuhnert 90] and *ATHENE II*, 3-wheeled vehicles, about 1.5 m long and 0.7 m wide, driven and steered by the front wheel. Main sensors are a video camera on a pivoting platform and odometers on the rear wheels. Four inbuilt behaviors, “going along a corridor”, “going around a corner”, “going towards a landmark”, and “stopping” enabled the robots to operate in a network of corridors.

Figure 2 (left) illustrates the guided exploration of an environment and the automatic generation of an attributed topological map. It shows the course traversed and (in a slightly abbreviated form) the instructions given by the guide at certain locations. The format of the instructions depends on the situation. At the start of the exploration the robot is informed of the name of the present location (“Rob-Lab”), the name of the next location (“Mani-Lab”), and the behavior to be executed (“Follow corridor”). The robot then follows the corridor and stops when the guide pushes a key to inform the robot that the destination has been reached. Then the robot receives the next set of instructions, and so on. Guiding the robot in the described way is easy, and it

takes hardly more time than just driving it along the same course by manual control.

Figure 2 (right) shows the resulting map, augmented by a later additional exploration covering the upper right part of the map. Although the map looks distorted, it is perfectly suitable for behavior-based navigation as the distortion reflects merely moderate, and thus irrelevant, errors of the geometric attributes, while the topology is correct.

3 Learning to Grasp Objects

General Approach

Similar to living organisms, future robots will not need any system calibration and no quantitative models of themselves. They will gather experiences through interaction with the world and continuously improve their performance based on the collected experiences. They will easily adapt to changes in the environment and to changes of their own system parameters.

A concept of a robot displaying some of those characteristics was introduced by [Graefe 1995]. It eliminates the need for a calibration of the robot and of the vision system, it uses no joint or world coordinates, and adapts automatically to changing parameters. The concept is based on a utilization of laws of projective geometry that always apply, regardless of camera characteristics, on a direct transition from sensor data to motion control commands, and on machine learning for the acquisition of knowledge regarding those system parameters which govern that transition.

The concept is simple. While the robot watches its end effector with two cameras, like a playing infant watches his hands with his eyes, it sends more or less arbitrary control commands to its motors. By observing the resulting changes in the camera images it "learns" the numerical values of a special Jacobian matrix, the Sensor-Control Jacobian [Graefe, Maryniak 1998], describing the relationships between such changes in the images and the control commands that caused them. After having executed a number of test motions in different configurations the robot is able to move its end effector to any position in the images that is physically reachable. If, in addition to the end effector, an object is visible in the images, the end effector can be brought to the object in both images and, thus, in the real world.

Once a robot based on this concept has localized objects in the images of its cameras, it can grasp them without any knowledge of its kinematics or its camera parameters. In contrast to other approaches with similar goals, but based on neural nets, no training is needed before the manipulation can start.

Initial Performance and its Improvement

The gripper approaches the object in a series of steps with some test motions in between. The grasping process is, therefore, relatively slow. In a first implementation the relationships between control commands and the resulting motions in the camera images were immediately forgotten after the end of each grasping experiment. This made the system insensitive to parameter changes between experiments, but it also prevented the robot from accumulating experiences for improving its motion skills.

Long-Term Memory

Ideally, a calibration-free robot should be able to start working immediately when it is switched on for the first time, without first requiring a training phase. Since it does not yet know its own characteristics, its initial movements will not be optimal, though. The robot should then learn from experience while it is performing its task, and improve its skill over time. Eventually the gripper should approach in a single direct motion any object to be grasped anywhere in the robot's work space without performing test motions or iterations. To do so, the robot must know the numerical values of the Sensor-Control Jacobians belonging to many different configurations of the arm covering the entire work space, and to different poses of the object. The robot should acquire this knowledge by learning and store the essence of what it has learned in some kind of long-term memory or knowledge base. Remembering is not enough, though. What has once been learned may later become inaccurate or incorrect, due to some gradual or sudden change of the robot's characteristics. It should then be modified or forgotten.

Self-Adapting, Forgetting and Relearning

In the real world the characteristics of the optical and mechanical subsystems of the robot may at any time change in an unforeseeable way, for instance due to the aging of parts of the robot or to some maintenance performed on it. In such a case the robot should be flexible enough to adapt the contents of its memory to the new situation. This may be done in one of two ways, depending on the nature and the magnitude of the changes. If the robot detects only minor differences between the anticipated results of motion control commands and the actually observed results, it applies a correction to some of the coefficients in the knowledge base; otherwise it continues to use the existing knowledge base. If the magnitude of the difference between anticipated and observed image motions exceeds a certain limit, the contents of the knowledge base are completely discarded and the robot begins to learn anew ([Nguyen, Graefe 2001]).

Structure of the Robot's Knowledge Base

Since the knowledge our robot uses consists mainly of numbers, namely the elements of Sensor-Control Jacobians, we use tables for storing this knowledge. In an initial implementation proposed by [Xie et al. 1997] the tables were multi-dimensional, where the dimensionality of the tables grew with the number of degrees of freedom (d.o.f.) that must be controlled. The approach worked well for a robot with 3 d.o.f., but the number of table elements and, thus, the necessary learning time would grow exponentially if the number of d.o.f. were increased. Therefore, in a more recent implementation we subdivide the knowledge base into a set of 1- or 2-dimensional tables, with each one of them containing knowledge relating to 1 or 2 specific joints of the robot ([Nguyen, Graefe 2001]).

Real-World Experiments

Setup. We have implemented the concept on an articulated arm robot (Mitsubishi Movemaster RV-M2) with 5 d.o.f., corresponding to the 5 joints J_1 to J_5 (Figures 3 and 4). The cameras have been attached to the robot on metal rods at

the first link so that they rotate around the (vertical) axis of Jt_1 together with the arm. They are mounted in a rather unstable way to make the impossibility of any calibration or precise adjustment obvious, and to allow easy random modifications of the camera arrangement.

Learning from Scratch. In a typical sequence of experiments the long-term memory was initially empty. An object was placed somewhere in the robot's workspace. Since no knowledge was available in this learning phase the robot moved the gripper slowly and in a sequence of steps towards the object and adjusted the orientation of the gripper according to the orientation of the object ([Graefe, Ta 1995], [Vollmann, Nguyen 1996]). After each step the control commands issued and the resulting image motions were recorded in the knowledge base.

In subsequent grasping operations, either the same object or different ones were used, and they were placed in positions or orientations different from the previous ones. If the object's pose was not too different from previously experienced ones, the robot could determine suitable control commands by interpolation, or extrapolation, based on the acquired knowledge and the current information from the vision system. The gripper was in such cases moved directly to the object in only one step.

Incremental Learning. An incremental learning ability was demonstrated by selecting the object in the learning phase to be a cylindrical one with a vertical axis of rotational symmetry. To grasp such an object, only joints Jt_1 , Jt_2 and Jt_3 need to be controlled. Hence, only the tables for these joints were filled with data. In a subsequent operation we used an object that does not have rotational symmetry around a vertical axis, forcing the robot to activate all of its 5 joints. Based on the acquired knowledge and using its first 3 joints it first moved the gripper to an *intermediate position* [Nguyen, Graefe 1997], [Graefe, Maryniak 1998]. Then it activated joints Jt_4 and Jt_5 to align the gripper step by step parallel to the object orientation. In the process it acquired control knowledge for Jt_4 and Jt_5 and added it to the -knowledge base to be used in later grasping operations.

Self-Adaption, Forgetting and Relearning.

To verify the self-adaption, forgetting and relearning abilities of the robot we arbitrarily modified the camera positions and orientations in the midst of experiments. In this case the robot continued to function, and, depending on the severity of the change of the poses (position and orientation) of the objects in the images, the previously acquired knowledge was either adapted, or a part of it, or all of it, was forgotten. Although the robot lost some of its skill when it discarded some or all of its knowledge, it re-acquired the necessary knowledge automatically in subsequent grasping operations. After a number of grasping operations had been performed, the knowledge base reflected the new situation, and the robot had regained its full skills.




Figure 3: Manipulator robot consisting of an articulated arm (5 degrees of freedom) with a two-finger gripper and a stereo vision system.




Figure 4: Schematic diagram of the arm and the camera arrangement.

4 Learning by Dialogue

Our latest example of a learning robot is the humanoid service robot *HERMES* that will eventually combine the above-mentioned navigation and manipulation skills [Bischoff 1997] (Figure 5). The robot is able to perform simple service tasks within initially unknown environments for users that may be initially unknown, too. Services handled by the robot include, but are not limited to, transportation tasks, entertainment, and gathering and providing information about people and their living or working environment. These services are based on several general robot skills that are combined and initialized depending on the robot's actual situation [Bischoff, Graefe 1999]. Obviously, communication between such a robotized "butler" and a human goes far beyond traditional man-machine interfaces.

Communicative Skills

In addition to a multitude of sensor, motor, sensorimotor and data processing skills, the robot possesses several communicative skills that are mostly based on natural language. Since natural language is the easiest and most natural mode of communication for a human it is desirable to integrate speech recognition and production into most personal and service robots. Language can be used to instruct the robot with higher-level goals or to intervene certain behaviors and modify their execution. Also, the ability to communicate by spoken and written natural language is an excellent basis for allowing a robot to learn from interaction with humans. However, to be accepted as cooperative partners, robots must not only have the ability to understand perfectly clear and complete utterances, but they must also resolve ambiguities and complement missing information – both being characteristics of human conversation. In doing so, an intelligent robot should pursue two approaches: One, it should use the current situation as a relevant context; two, it may evoke additional information and learn from the human through a dialogue.



Figure 5: Service robot *HERMES* used for human-robot interaction and cooperative learning experiments.

Situation-dependent context switching.

To enable both spoken and written natural language processing with limited computational resources, so-called contexts were designed

that cover all the commands, statements and questions that might be given to the robot. However, each context contains only those grammatical rules and word lists that are needed for a particular situation. When the robot knows what kind of answers it may expect from the user at a given moment it can switch to a context and disable or enable word lists as appropriate for the current situation. For example, when the robot asks for confirmation, whether it should execute a certain task or not, the answers will be most likely “yes” or “no” and it would not make any sense to expect, and to test, many other words.

Speech recognition. By limiting the set of recognizable words or phrases that can actually be expected, the risk of speech recognition mistakes, especially in environments with a high level of ambient noise, is reduced considerably. Another advantage of using situation-dependent context switching is that speech recognition can be speaker-independent if a proper engine is used. Our speech recognition engine enables speaker-independent recognition of continuous speech, which means that users may speak to the

system naturally, without pauses between words. This is a very important feature because it allows *anybody* to communicate with the robot without needing any training with the system.

Cooperative Learning

Two kinds of learning have been implemented with respect to the communicative skills of the robot. One, extending the lexical and syntactical knowledge bases of the robot. Two, introducing new knowledge to various databases in the course of *cooperative* learning of environmental features. The disambiguation of imprecise or unclear commands with respect to the prevailing situation is another aspect of learning that will be considered in the sequel.

The robot’s lexical and syntactical knowledge bases can be extended, firstly, by directly editing them (since they are text files), and secondly, by a dialogue conducted by the robot, thus adding new words, argument classes and prototypes. Also, new macro commands can be learned during run time.

Many types of dialogues exist to *cooperatively* teach the robot new knowledge and to build a common reference frame for subsequent execution of service tasks. For example, it is possible for the robot – besides learning an attributed topological map of the environment similarly to the *ATHENE* robots described in section 2 – to learn persons’ names, to learn how locations and objects are denominated by a specific person, where objects of personal and general interest are located, and how to grasp specific objects. This requires several databases and links between them in order to retrieve the required information, e.g., whose office is located where, what objects belong to specific persons and where to find them. To teach the robot names of persons, objects and places that are not yet in the database (and thus, cannot be understood by the speech recognition system), a spelling context has been defined that mainly consists of the international spelling alphabet. It has proved its effectiveness for our application. The dialogue sketched in Figure 7 illustrates the concept.

Obviously, there are some limitations in our current implementation. One limitation is that not all utterances are allowed or can be understood at any moment. The concept of contexts with limited grammar and vocabulary does not allow for a multitude of different utterances for the same topic. General speech recognition is not sufficiently advanced, and compromises have to be accepted in order to enhance the recognition in noisy environments. Furthermore, in our implementation it is currently not possible to track a speaker’s face, gestures or posture. This would definitely increase the versatility of the learning-by-dialogue concept.

5 Conclusions, Summary, Outlook

Equipping robots with learning and adaptation capabilities is necessary for making progress on the way to the realization of truly intelligent robots. Three robots capable of machine learning in various domains have been presented: a mobile robot, a stationary manipulator, and a service robot combining locomotive, navigational and manipulatory skills. All robots are endowed from the start with certain basic abilities, and are able to incrementally

Cooperative Learning Experiment
In this experiment two tables have been placed in the field of view of the robot. HERMES finds itself in a state of waiting for a new instruction within the (default) service context.

Person: “Hello!”
HERMES: “Hello. My name is HERMES. What is your name, please?”
Person: “My name is Rainer.”
HERMES cannot find the person’s name in its database, and therefore, switches to the spelling context.

HERMES: “Please spell your name!”
Person: “Romeo Alpha India November Echo Romeo.”
HERMES switches to the confirmation context.

HERMES: “Your name is Rainer. Is this correct?”
Person: “Yes!”
HERMES extends its database of persons and switches back to the service context.

HERMES: “Hello Rainer. Pleased to meet you. What can I do for you?”
Person: “Take this glass. Place it onto my table.”
These sentences invoke a number of different skills. First HERMES takes over the glass the person currently holds in his hand. Then, the robot resolves the pronouns “it” and “my” from its situated knowledge. However, the robot does not yet know which table the user has in mind. It could be either one of the two in its current field of view or a table at a different location. HERMES has to ask the user to help it solving the current ambiguity.

HERMES: “I have found two tables. Which one is yours?”
HERMES switches to the attributes context.

Person: “The left one!”
HERMES switches back to the service context, approaches the left table and places the glass onto it. Then, it extends its database of objects, and marks the approached table as one of Rainer’s tables.

Figure 7: Excerpt from a dialogue between a human and *HERMES* to place an object onto a table. In its course, *HERMES* learns more about its environment and stores this knowledge in several databases for later reference. It should be noted how often contexts are switched, depending on the robot’s expectations, thus improving the speech recognition considerably (Sentences written in italic describe the robot’s status or actions and context switching).

extend their knowledge and abilities through learning, as demonstrated in several real-world experiments.

The autonomous robots *ATHENE I* and *II* were used to study a special form of learning that may similarly be observed with animals exploring a new environment: the acquisition of knowledge about the topology and the geometry of a network of passageways. Experiments with a guided exploration of an environment showed that the system structure as implemented and the representation of knowledge in the form of an attributed topological map are indeed suitable for realizing machine learning. After an exploration run the robot was able to navigate autonomously in the explored environment.

The vision-guided stationary manipulator can successfully grasp objects of various shapes, although no knowledge regarding the optical or kinematic parameters of the robot is used. Even arbitrary rotations and shifts of the cameras are tolerated while the robot is operating. Key to this unusual robustness is the renunciation of quantitative model knowledge (i.e., camera model and kinematic model of the arm). Instead, the direct transition from image coordinates to motor control commands is learned, without using any inverse perspective or kinematic transformations.

The humanoid robot *HERMES* constitutes an excellent basis for investigating various aspects of machine learning and for demonstrating them in natural environments in close interaction with humans having had no special training. The robot is able to learn a user's wishes and many facts relating to the environment through dialogues in natural spoken language, even if the user expresses them in an ambiguous way, as humans often do when they speak naturally.

Presently, *HERMES'* learning ability is largely limited by its sensory abilities. Therefore, we plan to improve two of its sensory abilities, vision and touch.

If the robot had the ability to read door plates and other signs while traversing a corridor, it could more easily and more comprehensively gather navigational information while performing service tasks. If *HERMES* could understand gestures, as humans commonly use them when giving directions, this could make communication smoother and more robust.

HERMES has a touch-sensitive skin around its bumpers and on its fingers. By fusing touch data with vision data, the robot should be able to learn how to navigate safely in very narrow spaces and to manipulate many different kinds of objects dexterously. Covering the arms, too, with such a skin will open up the opportunity of having *HERMES* learn how to avoid collisions of its arms with objects in the environment and to touch objects gently with its arms to move them in a controlled way.

References

- Beer, R. (1990):** Intelligence as Adaptive Behavior. Academic Press.
- Bischoff, R. (1997).** *HERMES* – A Humanoid Mobile Manipulator for Service Tasks. Proc. of the International Conference on Field and Service Robotics. Canberra, Australia, Dec. 1997, pp. 508-515.
- Bischoff, R.; Graefe, V. (1999).** Integrating Vision, Touch and Natural Language in the Control of a Situation-Oriented Behavior-Based Humanoid Robot. IEEE Conf. on Systems, Man, and Cybernetics, Oct. 1999, pp. II-999 - II-1004.
- Graefe, V. (1995):** Object- and Behavior-oriented Stereo Vision for Robust and Adaptive Robot Control. International Symposium on Microsystems, Intelligent Materials, and Robots. Sendai, pp 560-563.
- Graefe, V.; Maryniak, A. (1998):** The Sensor-Control Jacobian as a Basis for Controlling Calibration-Free Robots. IEEE International Symposium on Industrial Electronics, ISIE '98. Pretoria, pp. 420-425.
- Graefe, V.; Ta, Q. (1995):** An Approach to Self-learning Manipulator Control Based on Vision. IMEKO Int. Symposium on Measurement and Control in Robotics, ISMCR '95. Smolenice, pp 409-414.
- Graefe, V.; Wershofen, K. P. (1991):** Robot Navigation and Environmental Modelling. IARP – Proceedings, Second Workshop on Multi-sensor Fusion and Environment Modelling. Oxford.
- Kuhnert, K.-D. (1990):** Dynamic Vision Guides the Autonomous Vehicle *ATHENE*. Japan-USA Symposium on Flexible Automation. Kyoto, pp 507-510.
- Nguyen, M.-C.; Graefe, V. (1997):** Object Manipulation Controlled by Uncalibrated Stereo Vision. The 2nd Chinese Congress on Intelligent Control and Intelligent Automation, CWCICIA '97. Xian, Vol. 1, pp. 77-83.
- Nguyen, M.-C.; Graefe, V. (2001):** Learning by Doing – an Approach to Robotic Skill Acquisition. SICE, Nagoya (CD Rom).
- Vollmann, K.; Nguyen, M.-C. (1996):** Manipulator Control by Calibration-Free Stereo Vision. In D. Casasent (ed.): Intelligent Robots and Computer Vision XV, Proceedings of the SPIE, Vol. 2904. Boston, pp. 218-226.
- Wershofen K.P., Graefe, V. (1992):** An Intelligent Autonomous Vehicle Guided by Behavior-based Navigation. IFToMM-jc International Symposium on Theory of Machines and Mechanisms. Nagoya, pp. 244-249.
- Xie, Q.; Graefe, V.; Vollmann, K. (1997):** Using a Knowledge Base in Manipulator Control by Calibration-Free Stereo Vision. IEEE Intern. Conference on Intelligent Processing System. Beijing, pp. 1307-1311.