Calibration-Free Robots for Cost-Effective, Dependable and Robust Automation

Volker Graefe

Intelligent Robots Laboratory
Bundeswehr University Munich, Germany

Graefe at unibw.de

Abstract—**Robots that never need any calibration of their kinematics, their actuators or their sensors promise great advantages in terms of maintenance cost, robustness and dependability. Here we propose an approach for realizing such robots for calibration-free vision-based object manipulation. The underlying concepts are based on the utilization of laws of projective geometry that always apply, regardless of camera characteristics, and in some cases on learning by doing. No quantitative models of the robot's kinematics, control characteristics or sensors are used. Consequently, even gross changes in those characteristics as they may be caused, e.g., by the aging of parts or by maintenance work, are tolerated, often without any degradation of the performance of the robot. The proposed approach is based on a vision system with an uncalibrated camera pair observing the robot's work space as the main sensor and a strategy that lets the robot learn automatically the relationships between motor control commands and resulting sensor data. Real robots controlled on the basis of the approach have proved their effectiveness, adaptability and robustness in extensive real-world tests.**

I. Introduction

Conventional robots depend heavily on quantitatively correct models of their own characteristics and of the relevant parts of their environments. For maintaining the necessary accuracy of the models they require repeated calibrations of their actuators and sensors, especially the vision system. This is a major impediment to the practical application of vision-guided robots in industrial, public, home and other less than perfectly structured environments. Not only does it constitute a significant cost factor, but a dependence on the accuracy of models and their parameters also prevents systems from behaving robustly in the real world where continuous and unforeseeable changes are the rule, rather than the exception. In contrast to conventional robots, animals and humans do not depend on any calibrations for controlling their motions. Especially humans can effortlessly adjust to changes in their sensory system, e.g., when putting on or off glasses, and in their environments, e.g., when driving different cars of unknown characteristics in unknown environments.

Robots depending neither on inbuilt quantitative models nor on pre-defined numerical values of any parameter and do not make any reference to external coordinate systems never need any explicit calibration. They promise great advantages in terms of robustness and cost of ownership. Such robots are called "calibration-free."

It may seem at first glance that methods not depending on world coordinates and on accurate internal models obtained by a careful calibration can yield only qualitative and inaccurate results. However, this is not correct. One counter-example is the motion stereo method introduced by [Graefe 1990] and [Huber, Graefe 1991]. It allows quite accurate distance measurements (relative errors less than 1 %) in the real world without any camera calibration. Moreover, nature provides an abundance of examples for systems that operate and survive in a wide variety of unstructured environments without ever needing an explicit calibration.

II. Approaches to Vision-Based Robot Control

2.1 The conventional, model-based approach

Conventional robot control depends on quantitatively accurate models of the robot, its sensors and actuators, and of the relevant objects in the environment. A stream of measurements from well-calibrated sensors is needed to continuously update the numerous variable parameters of the models.

As an example let us imagine an articulated arm robot, equipped with a stereo vision system that is supposed to grasp an object located somewhere in its work space. A classical approach for controlling such a robot would evaluate the camera images according to the well-known methods for stereo evaluation (a kind of inverse perspective transformation). This transformation requires all internal and external camera parameters (location, orientation, focal length, principal point, etc. of each camera) to be known with great accuracy. The result would be the location of the object to be grasped relative to some camera-fixed coordinate system. The coordinates of the object in a ground-based coordinate system would then be determined by an appropriate coordinate transformation again critically depending on a substantial number of parameters. Finally, the joints of the robot would be controlled in such a way as to move the gripper to that point in the ground-based coordinate system which was determined to coincide with the location of the object. This control implies another transformation (inverse kinematic transformation) between the ground-based coordinate system and the joint angles, which requires accurate knowledge of the robot's dimensions, kinematics, joint angles and control characteristics.

Even for this simple task the control of a vision-guided manipulator according to the traditional model-based approach depends critically on about 30 or 40 parameters. Their exact numerical values must be obtained by a hard-to-perform accurate calibration of the cameras and the arm's kinematics. Moreover, the control characteristics of all actuators must be known.

Model-based control breaks down when there is no accurate quantitative agreement between reality and the models. Discrepancies between reality and models may come about easily; an error in one of the many coefficients that are part of the numerical models may be sufficient for causing such a discrepancy and, thus, a failure of the robot. Among the many possible causes are initial calibration errors, aging of components, changes of environmental conditions such as temperature, humidity or illumination, maintenance work and replacement of components, to mention only a few. Consequently, robots depending on model-based control usually work only in carefully controlled environments and need frequent recalibrations, in addition to a cumbersome and expensive initial calibration.

2.2 The calibration-free approach

To overcome the shortcomings of model-based robot control many researchers have striven to develop robot control concepts that avoid, or at least reduce, the necessity of modeling and calibration. [Cooperstock, Milios 1993] use a set of neural networks to realize a robot that is able, after a training phase, to approach an object and grasp it without requiring a calibration. [Yoshimi, Allen 1994] use an uncalibrated camera system to perform a peg-in-hole task utilizing rotational invariance in a specific setup. [Hollinghurst, Cipolla 1994] avoid using an exact camera model by performing a self-calibration at four known positions. Doing this, they can handle calibration errors of the cameras as well as linear modeling errors of the manipulator.

A key point of the more general calibration-free approach introduced first by [Graefe 1995] is the way how the robot and the world are modeled (Figure 1). The robot controller models everything outside itself as one "black box" of unknown internal structure. The black box includes the entire external world, but also the sensors, actuators and other components of the robot. The robot controller can send control commands to the black box and receives sensor data from it. We humans know that control commands cause the robot's actuators to move and the state of the robot and of objects in the world to change, and that the sensor data reflect the state of objects in the world. The robot controller does not know any of these things. All it knows is that some sensor data will change in a more or less systematic, and thus predictable, way as a consequence of the control commands the controller issues, while other sensor data will change (actually, due to independent motions of external objects, but from the robot's point of view seemingly at random), even if the controller does not send any control commands. Instead of a world model a calibration-free robot uses a set of associations between control commands on one hand, and resulting sensor data, or changes in sensor data, on the other hand. It may acquire such knowledge either by being programmed accordingly, or by learning.

While a conventionally controlled robot is concerned with trajectory planning, world coordinates and accurate measurements, a calibration-free robot operates quite differently. To reach a specific goal it imagines what the sensor data should look like when the goal state, or some intermediate state, is reached. It then compares the set of actual sensor data with the set of imagined and desired sensor data and issues control commands to its actuators to minimize the difference between the two sets of sensor data. As a side effect the goal state in the real world is reached as soon as the difference between the two sensor data sets vanishes.

In the case of object-grasping, for example, the calibration-free robot knows implicitly that a suitable reference point that is fixed to the object to be grasped should coincide with the center of the robot's gripper before the gripper is closed. Coincidence of two points in the real world implies that they appear to coincide in both camera images. As long as that is not the case, the robot controller issues control commands that are computed from the distance between the two points *in the images*. The robot control makes no reference to any parameters describing the camera or the kinematics and dimensions of the robot or to any world coordinates. Also, the trajectory that the robot follows while it approaches the object is not computed; it simply emerges.

This approach eliminates the need for any numerical models of the robot and a calibration of its sensors. Because of its simplicity it may easily be augmented by a learning component for an automatic adaptation to changing parameters. It is based on the utilization of laws of projective geometry that always apply, regardless of camera characteristics, on a direct transition from visual input information to control commands (without any models of the robot and the external world, as indicated in Figure 1), and optionally on machine learning for the acquisition of knowledge concerning the relationships between control commands and subsequent changes in sensor data.

III. An Example: Object Manipulation

3.1. The task

The manipulation task studied for implementing and evaluating the calibration-free control concept is simple, but characteristic: An object that is placed at an arbitrary location somewhere in a robot's work space is to be grasped and picked up by the robot's 2-finger gripper. A manipulator arm (Mitsubishi Movemaster RV-M2) with 5 degrees of freedom, corre-

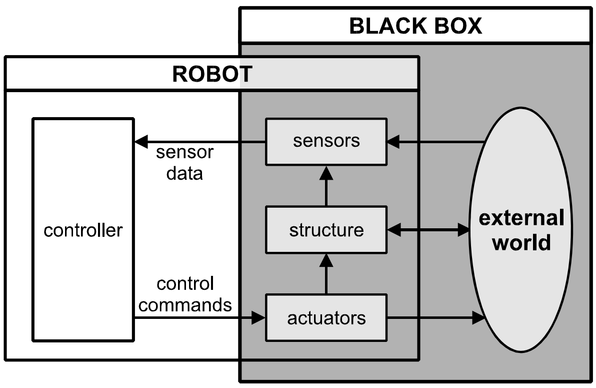

Figure 1: The controller of a calibration-free robot sees everything outside itself as one "black box", including the robot's structure, its actuators and sensors, and the world outside the robot.

sponding to its 5 joints, Jt_1 to Jt_5 (Figures 2 and 3), and equipped with a pair of video cameras may serve as an example of a calibration-free robot.

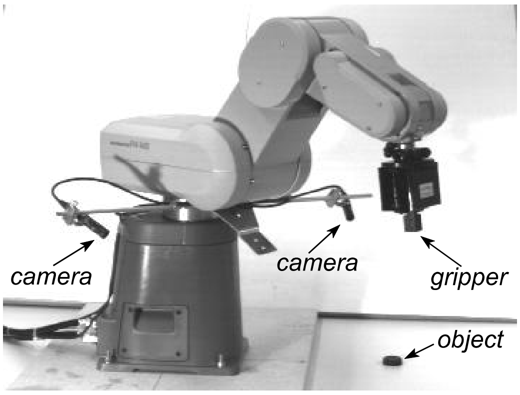

Figure 2: Articulated arm robot (5 degrees of freedom) with a two-finger gripper and two cameras as used for the experiments

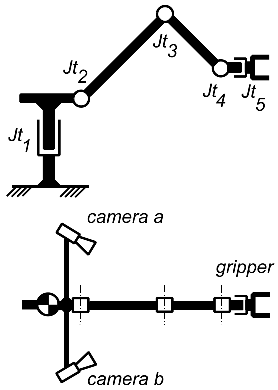

Figure 3: Schematic diagram of the arm and the camera arrangement.

The cameras have been attached to the robot on metal rods at the first link, so that they rotate around the (vertical) axis of the joint, 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 arbitrary modifications of the camera arrangement.

To simplify the explanation of the calibration-free approach, it will be assumed initially that the objects to be manipulated are of rotational symmetry (e.g., cylinders or spheres) with a vertical axis. Only 3 of the robot's 5 degrees of freedom (corresponding to Jt_1, Jt_2, and Jt_3) need then to be considered; Jt_4 may be controlled in such a way as to always keep the gripper in a vertical orientation, and the rotation of the gripper around its vertical axis (Jt_5) is irrelevant due to the rotational symmetry of the object.

In the context of this particular work, we are only interested in the problem of robot control without model knowledge, not in grasping strategies, collision avoidance or complex vision problems. Therefore, we have structured the environment for our experiments in such a way that both the gripper and the object are easy to locate in the images, and that the object is easy to grasp by a gripper approaching it from above without any risk of collisions with the table or other objects.

For the purpose of describing the concept of our approach we model both the gripper and the object as visible points. The robot's vision system recognizes and tracks the gripper and the object in the camera images, assigns reference points to them, and continuously delivers the image coordinates of both points in the images of both cameras (Figure 4). When the gripper has reached its grasping position the two points, gripper and object, coincide in the real world. Regardless of any camera characteristics, this is the case if, and only if, they also coincide in the images of both cameras. The rendezvous between the gripper and the object may, therefore, be brought about by any sequence of robot motions that makes the two points coincide *in the images* of both cameras simultaneously. This is always true, regardless of any specific characteristics of the robot or the cameras.

3.2 The control concept for objects of rotational symmetry

In implementing the concept we made the following assumptions:

- The robot's gripper and the object to be grasped are visible for both cameras; the vision system recognizes them and delivers continuously their image coordinates in both camera images.
- The internal and external camera parameters (optical characteristics, locations, and viewing directions) are unknown. Therefore, the functional dependencies of image coordinates on world coordinates are also unknown.
- The motions of the cameras are unknown. The cameras may either be at rest (located somewhere outside the robot), or they may move together with some part of the robot (not necessarily the end effector); they do not move independently, and they do not move relative to each other.
- The dimensions and joint angles of the robot are unknown; there is no kinematic redundancy.
- The robot's motions are controlled by a set, or vector, $\underline{c}$, of numerical control words (one control word for each motor) that is sent to the actuators. The functional dependencies of world coordinates of the gripper on control words are unknown since we do not use quantitative models. However, we assume that a zero control word causes no motion, while a control word of small magnitude causes a small motion, and a control word of larger magnitude causes a larger motion.

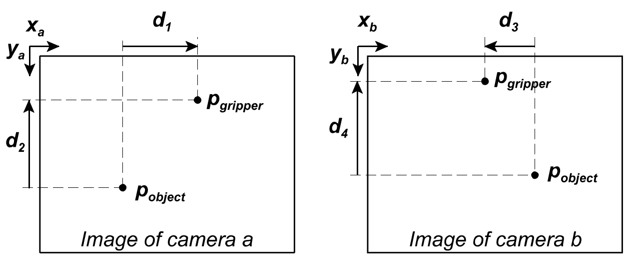

Figure 4: The gripper and the object (modeled as points) in the images of two cameras, a and b

If $\underline{d} = (d_1, d_2, d_3, d_4)^T$ (Figure 4) is the vector of image distances between the gripper and the object in all dimensions used (x and y in both images), then the task of making the gripper rendezvous with the object is equivalent to the task of making $\underline{d}$ vanish.

To solve this task without knowing the robot's kinematics or control characteristics we define a special Jacobian matrix, J, the Sensor-Control Jacobian. It describes relationships between control words (c_n) that have been transmitted to the actuators, and resulting changes in sensor data (here: changes in image distances). In the case of our robot where each control word (when sent to its actuator) causes a proportional change in the corresponding joint angle, J has the form given in (1):

$$J = \begin{bmatrix} \frac{\Delta d_1}{c_1} & \frac{\Delta d_1}{c_2} & \frac{\Delta d_1}{c_3} \\ \frac{\Delta d_2}{c_1} & \frac{\Delta d_2}{c_2} & \frac{\Delta d_2}{c_3} \\ \frac{\Delta d_3}{c_1} & \frac{\Delta d_3}{c_2} & \frac{\Delta d_3}{c_3} \\ \frac{\Delta d_4}{c_1} & \frac{\Delta d_4}{c_2} & \frac{\Delta d_4}{c_3} \end{bmatrix} \qquad (1)$$

If $\underline{d}$ is the vector of image distances between the gripper and the object, and J is the Sensor-Control Jacobian matrix, issuing a control word vector, $\underline{c}°$, satisfying

$$-\underline{d} = J \cdot \underline{c}° \tag{2}$$

would make $\underline{d}$ equal to $\underline{0}$, if the system were linear. Because of the nonlinearity of the (unknown) function between control commands and image distance, outputting $\underline{c}°$ will generally not make $\underline{d}$ become zero. However, it will decrease its magnitude, and iterating the process will finally lead to a sufficiently small image distance vector, $\underline{d}$, (and a correspondingly small distance in the real world) such that the grasping task may be accomplished (for more details see [Graefe, Maryniak 1998]).

It should be noted that (1) and, thus, (2) contain only such quantities that are directly available to the robot's controller without requiring the use of any coefficients that would have to be determined by some calibration procedure. Specifically, the image distances, d_n, are raw data as acquired by the vision system and they are measured in arbitrary units internal to the vision system. The magnitude of those units is irrelevant since multiplying all image distances, d_n, by the same nonzero constant, k, (i.e., some calibration factor) would not affect $\underline{c}°$ as k would then appear on both sides of (2). The control words, c_n, are dimension-less numbers generated by the robot controller itself. A control algorithm that uses exclusively quantities that are internal to the controller or the sensor system is the essence of the calibration-free approach.

One problem with equation (2) is that there are 4 sensor data and, thus, 4 equations for computing the 3 control words forming the control word vector, $\underline{c}°$. Consequently, the system of equations, (2), is over-determined. This problem may be addressed in various ways, e.g., by omitting one of the equations [Graefe, Maryniak 1998], by a transformation from the original 4-D sensor data space into a 3-D data space [Maryniak, Graefe 1998], or by utilizing the redundancy inherent in the over-determined system for improving the robustness of the robot control [Maryniak, Graefe 1999].

For any given robot the Sensor-Control Jacobian, J, depends on the situation, i.e., the actual location of both the gripper and the object. A simple way to determine it in any given situation is to let the robot execute test motions, one at a time, for each degree of freedom, and to measure the resulting change of the image distance vector, $\underline{d}$. The test motions should be executed at the beginning of each grasping process and then again at the beginning of each iteration step. This method works very well and it lets the robot tolerate even gross changes of its camera characteristics. The disadvantage of this method is that the numerous test motions make the grasping slow – too slow for most applications.

We have studied two ways of speeding up the operation of the calibration-free robot. The first one avoids test motions by determining the Sensor-Control Jacobian "on the fly", i.e., during the normal motions of the robot. It is based on an approach proposed by [Jägersand et al. 1997). We have implemented this method and tested it successfully, but since the mathematics behind it are too complex to fit within the frame of this paper we refer to [Maryniak 2002] for details.

The second way for speeding up the grasping operation involves a kind of machine learning where the robot remembers the Sensor-Control Jacobian that it has measured in one situation and recalls it when it later encounters a similar situation again. It will be described in section 3.4.

3.3 Grasping objects of general shape and orientation

Figure 5 shows examples of objects that lack rotational symmetry, but have such a shape and orientation that they may be grasped from above. To grasp such objects when they are lying on a flat table, a fourth degree of freedom of the robot, the one associated with joint Jt_5 (Figure 3), must be activated. To grasp objects as shown in Figure 5 in all possible spatial orientations, an arm with at least 6 degrees of freedom is necessary. If, as it is the case with our robot, only 5 degrees of freedom are available, the objects can be grasped in many, but not all spatial orientations.

The method we have developed for this is an extension of the previous one. First, by controlling joints Jt_1, Jt_2, and Jt_3 (Figure 3), the gripper is brought to a location where the image of its center point is near the image of the object reference point in both images (e.g., about 1 diameter of the gripper away from it). Then the gripper is aligned with the main axis of the object by applying a similar strategy as before, only that now the joints, Jt_4 and Jt_5, are controlled in such a way that the angles between one of the gripper's edges and the main axis of the object *in the images of both cameras* are minimized. Generally, when two lines in an image, such as the edges of the gripper, are parallel to each other, it does not imply that the corresponding lines in the world are parallel, too. However, when the gripper is close to the object, as it is in the final phases of the grasping process, this error is small enough to be neglected. For details see [Nguyen, Graefe 1997].

Due to the kinematics of the articulated arm, moving joint Jt_4 changes both the orientation and the location of the gripper. Therefore, the approach and alignment of the gripper to the object must be iterated until both the orientation and the location of the gripper are simultaneously approximately correct.

3.4 Learning and Self-adaptation

As mentioned in section 3.3 above the execution of test motions makes the operation of the robot too slow for many practical purposes. To remedy this we have developed a strategy where the robot remembers the accumulated control words that have made it assume specific configurations and enabled it to

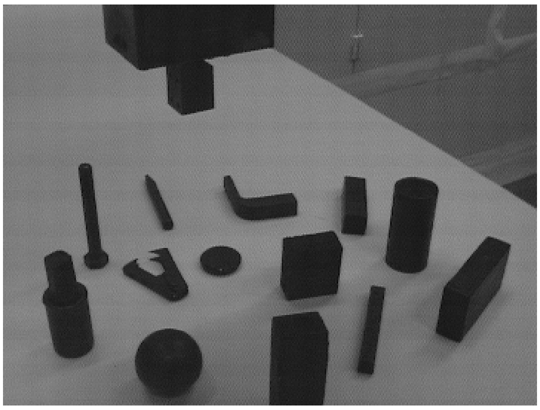

Figure 5: The gripper of the robot and some objects that were used in our experiments, as seen by one of the robot's cameras

grasp objects of specific poses (apparent locations and orientations in the camera images). When the calibration-free robot is to grasp an object for the first time, it has no knowledge to support the operation, and it must execute test motions. When it is to grasp another object of the same pose, it simply executes the remembered control command and approaches the object in one single step. It may then be said that the robot has learned how to grasp objects of that particular pose fast and efficiently.

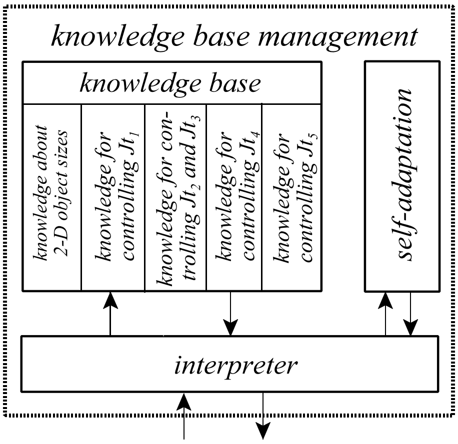

Figure 6: The knowledge module is structured into a set of independent sub-modules.

When the robot is to grasp for the first time an object at a second pose that it has not experienced before, it will again execute test motions. After that it can efficiently grasp objects at both poses and, by interpolation and extrapolation, also at intermediate and nearby poses. After having grasped objects of various poses distributed over its work space once, the robot can grasp objects anywhere in its work space directly and without performing any test motions.

Ideally, an uncalibrated robot should be able to start working immediately after it has been switched on, without requiring a training phase. Since it does not yet know its own characteristics, its initial movements may not be optimal, though. The robot should then learn from experience while it is performing its tasks and improve its skill over time. This is exactly how a robot employing the above control strategy behaves. However, the characteristics of the robot or of the environment may change, for instance, due to the aging of parts or to some maintenance that is performed on the robot. In such a case the robot should be flexible enough to modify gradually what it has learned or to forget it completely, depending on the nature of the changes, in order to adapt to the new situation.

In principle it would be possible to use a neural network for implementing the learning function. However, a neural network needs an often lengthy training before it can be used. Moreover, it is difficult to control what is actually being learned by a neural network, and incremental learning and gradual forgetting are also hard to realize. Therefore, we have implemented the knowledge base of our robot as a set of 1-D and 2-D tables where each table contains those elements of the Sensor-Control Jacobian and their derivatives that are needed for controlling one or two of the robot's motors. An additional table contains data supporting the recognition of the gripper and of objects to be grasped. Although it would in principle be possible to use a single table of higher dimensions, the size of of the table would grow exponentially with the number of degrees of freedom that are to be controlled. For more than 3 degrees of freedom such a table would be very large and difficult to maintain. Figure 6 illustrates the structure of the knowledge module as we have realized it for our robot. It contains not only the knowledge base itself, but also other separate submodules (for details see [Nguyen, Graefe 2001]).

The self-adaption submodule is responsible for controlling the adaptation of the available knowledge to changes of the robot's parameters or of the environment. After each motion of the robot it computes the motion error, i.e. the difference between the actual image distance vector, $\underline{d}$ and the previously expected image distance vector, $\underline{d}$*. If the magnitude of the motion error is small (below a certain threshold), it is assumed that it is caused by unavoidable noise, and it is neglected. If the magnitude of the motion error is larger, but still moderate (below a second threshold), those data in the knowledge base which have contributed to the motion error are modified accordingly. This modification does not slow the robot down since it does not involve any additional motion. If the magnitude of the motion error is very large and exceeds the second threshold, it is assumed that the characteristics of the robot or its sensors have grossly changed and the previously acquired knowledge is no longer valid. In such a case the knowledge base is erased completely and the robot begins to learn from scratch again.

It is remarkable that the learning, adaptation, forgetting and relearning occur without any human intervention. Moreover, the robot always continues to work, even while it is learning or relearning, although its working speed is temporarily reduced after the knowledge base has been erased.

IV. Real-World Experiments

A series of experiments was conducted where the robot shown in Figure 1 grasped various objects, some of which are shown in Figure 5. The objects were placed at arbitrary locations and orientations in the robot's work space; certain orientations that would have required an arm with 6 degrees of freedom were avoided. [Video Grasp] shows some of the experiments. The video also shows how changes in the orientation of a camera were tolerated without any deterioration of the robot's performance since it only led to a quick modification of the robot's knowledge base. A major change in the camera's location was also tolerated, but it caused the robot to erase its knowledge base and relearn its grasping skill from scratch which made it slow down temporarily.

Other experiments were conducted with a humanoid robot that implemented the concepts introduced here not only for object manipulation, but also for navigation. Some videos showing it in operation may be found at [Videos *HERMES*]. That robot absolved a long-term test when it operated in a museum for 6 months. It worked for several hours every day and functioned correctly despite the unavoidable aging of parts and random changes of system parameters.

V. Summary and Conclusions

The concept of calibration-free robots and ways for implementing it were introduced. In contrast to conventional robots which for their control depend critically on substantial numbers of hard-to-determine parameters and on operations such as inverse perspective or inverse kinematic transforms, calibration-free vision-guided robots require no quantitative knowledge of the parameters of their sensors and their kinematics. Instead, they use direct transformations between sensor data (e.g., image coordinates) and motor control words. Based on this concept a robot can grasp objects 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 may start.

By renouncing the utilization of world coordinates and by not attempting to quantitatively model the robot, the world, and the interactions between the robot and the external world, calibration-free robots achieve a high degree of robustness and adaptability. This was demonstrated by arbitrarily changing the camera configuration while the robot was working. Also, in a long-term test one of our robots operated daily in a museum for 6 months and tolerated changes of its characteristics caused by maintenance and aging of parts. This robustness provided by the calibration-free approach will be especially useful for future service and personal robots that, unlike today's industrial robots, must cooperate with untrained humans and work in unstructured environments without having access to frequent maintenance and professional support.

References

Cooperstock, J. R.; Milios, E. E. (1993): Self-supervised learning for docking and target reaching. Robotics and Autonomous Systems. Vol. 11, pp. 243-260.

Graefe, V. (1990): Precise Range Measurement by Monocular Stereo Vision. Japan-USA Symposium on Flexible Automation. Kyoto, pp. 1321-1324.

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.

Hollinghurst, N; Cipolla, R. (1994): Uncalibrated Stereo Hand-Eye Coordination. Image and Vision Computing. Vol. 12, No. 3, pp 187-192.

Huber, J.; Graefe, V. (1991): Quantitative Interpretation of Image Velocities in Real Time. IEEE Workshop on Visual Motion. Princeton, pp 211-216.

Jägersand, M.; Fuentes, O.; Nelson, R. (1997): Experimental Evaluation of Uncalibrated Visual Servoing for Precision Manipulation. IEEE International Conference on Robotics and Automation. Albuquerque, pp. 2874-2880.

Maryniak, A. (2002): Vision-Based Manipulator Control of Calibration-Free Robots. Dissertation, Faculty of Aerospace Engineering, Bundeswehr University Munich.

Maryniak, A.; Graefe, V. (1998): Transforming Sensor Data to Increase Robustness in the Control of Calibration-Free Robots. IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS '98. Victoria, pp. 1834-1839.

Maryniak, A.; Graefe, V. (1999): Utilizing Sensor Data Redundancy to Gain Robustness in the Control of Calibration-Free Robots. IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS '99. Kyongju, pp. 519-524.

Nguyen, M.-C.; Graefe, V. (1997): Object Manipulation Controlled by Uncalibrated Stereo Vision. The Second Chinese Congress on Intelligent Control and Intelligent Automation, CWCICIA'97. Xian, Vol. 1, pp. 77-83.

Nguyen, M.-C.; Graefe, V. (2001): Object Manipulation by Learning Stereo Vision-Based Robots. IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS '01. Maui, pp. 146-151.

Video Grasp http://www.unibw.de/robotics/videos/

Videos *HERMES* http://www.unibw.de/robotics/videos/

Yoshimi, B. H.; Allen, P. K. (1994): Active, uncalibrated visual servoing. IEEE International Conference on Robotics and Automation. San Diego, Vol. 4, pp. 156-161.