ICIPS, Beijing, October 1997

Using a Knowledge Base in Manipulator Control by Calibration-Free Stereo Vision

Qiang Xie, Volker Graefe, Karl Vollmann

Institute of Measurement Science
Federal Armed Forces University Munich
85577 Neubiberg, Germany
Phone: +49 89 6004 3343
Fax: +49 89 6004 3074
e-mail: {qiang.xie|graefe|karl.vollmann}@UniBw-muenchen.de

ABSTRACT

A 4-dimensional knowledge base that may be used as the long-term memory for object manipulation by a completely uncalibrated vision-guided robot is introduced. It is constructed automatically by the robot during its normal operation and adapts itself to changing conditions. Thus, the system presents self-learning characteristics. The know-ledge base has been successfully tested in real-world experiments involving the grasping of objects.

1. INTRODUCTION

Conventional stereo vision methods for grasping objects require repeated calibration of the manipulator and the vision system. To avoid the protracted and, thus, expensive calibration procedure an approach to robust, adaptive and calibration-free manipulator control has been proposed by [Graefe, Ta 1995]. This method was later named “object- and behavior-oriented stereo vision” [Graefe 1995]. As an extension of this method [Vollmann, Nguyen 1996] have achieved manipulation of elongate objects in an arbitrary orientation, in addition to flat cylindrical objects. The great advantage of this trial-and-error-based method is that it requires no pre-knowledge. The disadvantages are that it is relatively slow due to repeated test motions, and that even after repeated experiments no experience is built up to improve the performance, since no long-term memory was included in the system.

This paper describes an approach for overcoming these shortcomings. Knowledge regarding the relationship between control commands and their resulting effects as observed in the camera images is collected and stored in a knowledge base whenever the robot moves. After having executed a sufficient number of movements and, thus, explored its work space the robot is able to approach any desired location in the work space more or less directly and much faster than by the original trial-and-error-based method.

However, accumulating knowledge only once is not enough to make the robot robust. Many factors may potentially affect the robot, the vision system, or the environment, and modify relevant system parameters either gradually or suddenly. If this happens the stored knowledge will no longer be correct and must be modified, or even discarded.

Therefore, in addition to the ability of learning, the robot also needs the capacities of forgetting and relearning. If a robot has such abilities it may begin working with an empty knowledge base and gradually fill it with data as it executes motions. As more and more data become available the robot will operate more and more skillfully. If some system parameters happen to change, the acquired knowledge will no longer be completely valid. The robot will then sometimes execute erroneous motions and correct them by the trial-and-error-method. Whenever this happens the robot adapts its knowledge base to the new situation and eventually regains its full level of skill.

In the sequel we will describe a system that enables a robot to learn and relearn. It extends the method of object- and behavior-oriented stereo vision for manipulator control [Graefe 1995] by including a 4-dimensional knowledge base that is used as a long-term memory. Before describing the new method in detail we will briefly introduce the experimental setup and the manipulation task that we have used for developing and testing the concept.

2. EXPERIMENTAL SETUP

2.1 The Robot

A specific setup has been used in real-word experiments designed to test our approach of using a knowledge base for vision-based manipulator control. A 5-degree-of-freedom articulated arm (Mitsubishi Movemaster 2) equipped with two video cameras is used (Figure 1). The two cameras participate in the rotation of the arm around its vertical axis, but they are fixed relative to the work plane of the robot. The location and orientation of each camera are somewhat arbitrary and not exactly known, but each camera should be mounted in such a way that its field of view covers that area within the work plane in which objects are to be manipulated. Of the 5 degrees of freedom of the robot arm, one refers to the rotation of the gripper around its axis. The remaining 4 degrees of freedom correspond to the joints J_0 ... J_3 (cf. Figure 2). Joints J_1, J_2 and J_3 allow the arm to be moved within a certain section of the vertical plane, the work plane. Joint J_0 allows the arm to be rotated around a vertical axis, thus determining the azimuth of the robot's work plane. In our experiments joint J_3 was always controlled in such a way that the gripper was in a vertical orientation, as indicated in Figure 2. Our arm has, thus, three independent degrees of freedom remaining, corresponding to the three joints, J_0, J_1 and J_2.

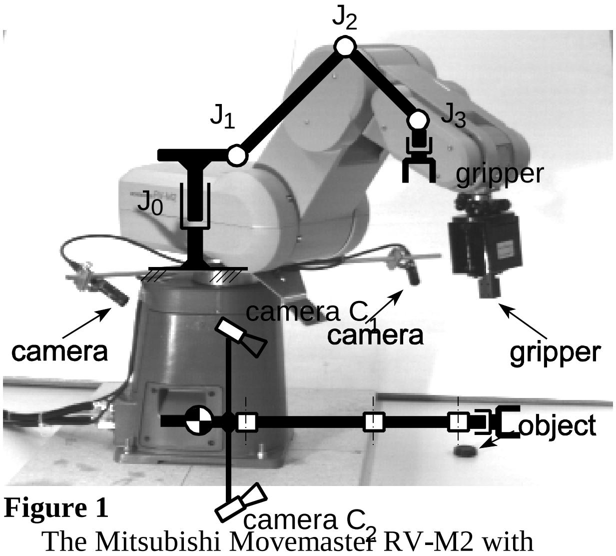

Figure 1
The Mitsubishi Movemaster RV-M2 with mounted cameras

Figure 2
The robot arm joints and the camera arrangements

2.2 The Task

Suppose, a (point-like) object that is to be grasped is visible in the images of both cameras. In order to pick up the object the (center point of the) gripper must be moved towards the object. If we knew the joint angle control words that correspond to the location of the object the task would be trivial.

However, since our robot and its vision system are uncali brated we do not have such knowledge. The task is then a twofold one:

- to bring the gripper to the object without needing additional information (e.g. according to the method of [Graefe 1995]) and
- to learn something in the process which enables the robot to execute a similar operation more skillfully in the future.

3. THE KNOWLEDGE BASE

3.1 The Concept

Each point in that part of the real world that is visible to one of the robot's cameras is transformed into the camera's 2-D image space by a perspective transformation. We may conceptually combine the image spaces of the two cameras into one 4-D image space. Any point in the 3-D object that is visible to both cameras is transformed into a point of the 4-D image space (it should be noted, however, that not all points in the 4-D space have a corresponding point in the real object space). As

Figure 2 shows, the cameras are attached to the robot in such a way that they follow its rotations around its vertical axis (joint J_0). Therefore, the absolute azimuth angle of an object point cannot be determined by image interpretation; only the azimuth angle relative to the current orientation of the work plane of the robot can be obtained.

The configuration of the robot is determined by three control words, C_0, C_1, C_2 that correspond to the angles of the joints J_0, J_1, J_2. The coordinates of any point in the work space of the robot may be described by that control word vector that would make the gripper center point coincide with the point in question. Since the absolute azimuth angle of an object point is not observable by our sensors we are only interested in object coordinates relative to the work plane of the robot. They are given by C_0', C_1, C_2, where C_0' stands for that increment in the control word C_0 that will cause the work plane to rotate by such an angle out of its present orientation that it coincides with the object point. The variables, C_0', C_1, C_2, generate a 3-D space, the modified control word space, and each object point in the work space is at any given moment transformed by the inverse kinematics of the robot into a point in the modified control word space. If the point is visible by both cameras it is also transformed into a point of the 4-D image space by the two cameras with their perspective transformations. There is, thus, a one-to-one correspondence between a section of the 3-D modified control word space and a section of the 4-D image space, but with our uncalibrated system we do not know the parameters of the correspondence quantitatively.

We may build up quantitative knowledge concerning the correspondence by experiments of the following type. We bring an object into the work space of the robot where it can be seen by both cameras. By image interpretation we obtain the coordinates of the object in the 4-D image space. Then we make the gripper meet the object (for instance, by the method of [Graefe 1995]) and, thus, obtain the object's coordinates in the modified control word space. By repeating this experiment for a number of object locations we may generate a transformation table containing for each one of those points both its 4-D image coordinates and the corresponding control word vector.

In fact, we do not even have to perform such experiments explicitly. After all, the described "experiments" are exactly the kind of operations that the robot performs as part of its normal activity. Suppose that the robot sees a point-like target in its work space. By image interpretation it obtains the object coordinates in the image space. If the transformation table contains an entry for that particular point in the 4-D image space the control word vector belonging to that point may be read from the table, and the gripper may be steered directly to the target. In the more likely case that no entry exists in the table for the exact target position we may try to find one or more points in a 4-D vicinity of the image of the target for which data already exist, and by a suitable interpolation move the gripper to a point near the target.

To actually reach the target point there are two alternatives: if the derivatives of the transformation from image coordinates to control words are recorded in the knowledge base, e.g., in the form of a Jacobian, and if the remaining distance to the target is not too long, the target may be directly approached. If the derivatives are unknown, their values may be obtained by executing test motions for each degree of freedom of the robot. Once the gripper has reached the target we may determine its coordinates in the modified control word space and update the transformation table by including an additional entry or by modifying entries belonging to neighboring points.

The robot, thus, acquires its knowledge in a form of learning by doing.

3.2 Practical Considerations

We have decided to base the transformation table on a regular grid in the image space with a grid constant of 64 pixels in each dimension. Since our images comprise 512 pixels * 512 pixels each there are 9^4 = 6.561 grid points (also called supporting points). The table allows this same number of entries, although some of the grid points do not correspond to points in the real 3-D world.

Both, the choice of a regular grid and the value of the grid constant are arbitrary, but we have considered them to be good enough for testing the concept in a first implementation. A dilemma exists when deciding on the right size for the grid constant. If it is too large, the nearest grid point will often be rather far away from the target point, and the error in approaching the target may be significant. On the other hand, if the grid constant is too small the resulting table is very large, and it takes many motions of the robot to fill a substantial portion of it with data, which is a prerequisite for being able to observe a significant learning effect.

3.3 The Contents of the Knowledge Base

According to our representation there are 6.561 supporting points in the 4-D knowledge base, each of which is defined by four coordinates as follows:

$$\begin{aligned} &P_{S_n} = (64*i,\ 64*j,\ 64*k,\ 64*l). \\ &n = l+9(k+9(j+9*i))=0,1,...,6560; \\ &i,\ j,\ k,\ l = 0,1,...,8. \end{aligned} \qquad \textbf{(1)}$$

A control word vector may be stored for each supporting point:

$$\boldsymbol{C}_{S_i} = (C_0',\ C_1,\ C_2)_{S_i}\ , \qquad \textbf{(2)}$$

where C_0', C_1, C_2 control the joint angles, α_0, α_1 and α_2, corresponding to the joints, J_0, J_1 and J_2, respectively. This control word vector is needed to make the gripper center point coincide with that point in the 3-D object space which corresponds to this supporting point.

4. APPLICATION OF THE KNOWLEDGE BASE

4.1 Knowledge Recall

When an object is to be grasped a target point near the object is determined first. The target point should be a certain distance away from the object to allow the robot to approach the target point without any danger of colliding with the object, but near enough to allow a straight forward grasping operation once the target point has been reached. The position of the target point in the 4-dimensional image space is determined by the image coordinates of the object as given by both cameras. According to the position of the target point the control word vectors stored in the surrounding 16 supporting points are read out; they form a set of control word vectors:

$$SET(\boldsymbol{C}) = \{\boldsymbol{C}_{S_i} | i=0,1,...,15\}. \qquad \textbf{(3)}$$

In order to obtain the control word vector corresponding to the target point from the above vector set, the Euclidean distance between the target point (P) and each of the surrounding 16 supporting points (P_{si}) in the 4-dimensional space is used to form the weight for each control word vector in the set:

$$w(\boldsymbol{C})_{S_i} = 1 - \frac{d_e(P,\ P_{s_i})}{d_{eMAX}} \qquad i=0,\ 1,\ ...,\ 15, \qquad \textbf{(4)}$$

here d_{eMAX} is the maximal distance between two supporting points; in our application with a grid constant of 64 pixels d_{eMAX} =128.

With those weights and using the method of linear interpolation the control word vector for moving the gripper to the target point is obtained by

$$\boldsymbol{C} = \frac{\sum_{i=0}^{15} w(\boldsymbol{C}_{S_i}) * \boldsymbol{C}_{S_i}}{\sum_{i=0}^{15} w(\boldsymbol{C}_{S_i})}. \qquad \textbf{(5)}$$

if the resulting position of the gripper is not yet close enough to the desired position, then the original trial-and-error-based method from [Graefe 1995] is adopted for a final approach to the target point.

4.2 Relearning

If the knowledge-based motion alone did not bring the gripper sufficiently close to the target, then the data in the knowledge base, i.e., the set of the control word vector that has been used, should be updated. Our strategy for doing this is to use the weights obtained from (4) once again with the following modifying algorithm:

$$\boldsymbol{C}_{S_i}^{\ new} = \boldsymbol{C}_{S_i}^{\ old} + w(\boldsymbol{C})_{S_i} * (\boldsymbol{C} - \boldsymbol{C}_{S_i}^{\ old}). \qquad \textbf{(6)}$$

4.3 Initialization of the Knowledge Base

For initializating the knowledge base a very simple method is used: if the gripper has been moved to a target point, and if some (or all) of the 16 surrounding grid points do not yet have valid data entries, then the control word vector that has been used moving the gripper is assigned to those grid points.

5. IMPLEMENTATION AND RESULTS

The described method was implemented and tested in real-world experiments. Two single-board image processors using the TMS 320C40 signal processor (one for each camera) were employed for locating and tracking the gripper and the target object in the images. They are physically part of a PC that performed all other knowledge processing and calculation.

When the knowledge base is empty (before the learning has started) or when a point is to be approached for which no data in the knowledge base exist, the trial-and-error-based method described in [Graefe 1995] is used. In this case grasping an object requires about 50 s to 60 s. If the object to be grasped is near a location that has already been reached by the robot in a previous experiment then a control word vector corresponding to that location may be recalled from the knowledge base, and the object is approached directly. In such a case the entire grasping process requires only 10 s - 20 s.

As more and more experiments are performed, more and more data are accumulated in the knowledge base, and more and more often a fast grasping is achieved. This is a form of skill acquisition by machine learning.

6. CONCLUSIONS AND DISCUSSION

A method that allows a completely uncalibrated vision-guided robot to manipulate objects has been introduced. It is based on a direct transition from camera image coordinates to motor control commands and bypasses the otherwise necessary inverse perspective and kinematic transformations. Compared to previously introduced similar methods, e.g., [Cooperstock, Milios 1993] or [Ritter et. al 1991] it has the additional advantage that the robot not only learns autonomously how to perform the task of grasping an object, but also remembers what it has learned and, thus, gradually builds up its skill.

In the real world the characteristics of the optical and mechanical subsystems of the robot may at any time change in an unforeseeable way. The proposed method includes the ability of adapting to such changes by relearning. This may be demonstrated, for instance, by arbitrarily modifying a camera orientation while the robot is working. In such a case the robot continues to function, although it temporarily looses part of its previously acquired skill and, therefore, moves more slowly. After a number of grasping operations have then been performed the knowledge base will have updated itself automatically to reflect the new situation, and the robot will have regained its full skill and operating speed.

Despite these successes there is still room for improvements, and some questions are still open. We hope to address the following aspects in the near future:

- Interpolation between neighboring supporting points may not be the best possible method for generating data from the knowledge base. Using only one support point and the local Jacobian to compute the correct motor control commands may be a better approach, although generating the Jacobian may not be an easy task.
- A more sophisticated form of relearning is desirable to accommodate both slow, gradual changes and also sudden, large changes. In the first case an approach based on the Kalman filter may be best, while in the second case a complete restart with an empty knowledge base may be better. However, it is not clear yet how the robot can determine the nature of changes that may have occurred.
- The decision to use a fixed number of supporting points located on a regular grid in the image space was a rather arbitrary one. It may be better to let the robot determine the number of supporting points and their locations automatically in an optimal way. The dilemma is that if the number of supporting points is too small it will often happen that no supporting point exists that is sufficiently close to the actual target point, and if the number is too large it will take unreasonably long to collect the data for all those points. (At least one movement of the robot is necessary to obtain the data for one supporting point.)
- Using a 4-D knowledge base, while straight-forward when 2 images are used, is probably not optimal. It leads to a relatively large number of supporting points, and some of the points on the regular supporting point grid do not correspond to any point in the 3-D object space. We are currently studying methods of using a combination of two 2-D and one 1-D knowledge base, or a single 3-D one.

We are convinced that the approach presented here will prove especially valuable for mobile and service robots operating in environments that are much less structured and predictable than typical factories where nowadays robots are mostly employed.

REFERENCES

Cooperstock, J. R.; Milios, E. E. (1993): Self-Supervised Learning for Docking and Target Reaching. Robotics and Autonomous Systems 11 (1993), pp 243-260

Graefe, V.; Ta, Q.-H. (1995): An Approach to Self-Learning Manipulator Control Based on Vision. IMEKO International Symposium on Measurement and Control in Robotics. ISMCR'95. Smolenice, pp 409-414

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

Ritter, H.; Martinez, T.; Schulten, K. (1991): Neuronale Netze: Eine Einführung in die Neuroinformatik selbst-organisierender Netzwerke. Addison Wesley, Bonn, München (u.a.), pp 179-199

Vollmann, K.; Nguyen, M.-C. (1996): Manipulator Control by Calibration-Free Vision. In D. Casasent (ed.): Intelligent Robots and Computer Vision XV, Proceedings of the SPIE, Vol . 2904. Boston, pp 218-226