MVA2000 IAPR Workshop on Machine Vision Applications, Nov. 28-30, 2000, The University of Tokyo, Japan

3—36 Grasping Objects of Various Shapes by Calibration-Free Vision-Guided Robots

Volker Graefe and Minh-Chinh Nguyen

Bundeswehr University Munich
The Institute of Measurement Science
85577 Neubiberg, Germany
Phone: +49 89 6004-3590, fax: +49 89 6004-3074
URL: http://WWW.UniBw-Muenchen.de/campus/LRT6
E-mail: Graefe@UniBw-Muenchen.de

Abstract

An approach to vision-based control of calibration-free robots is introduced. Similarly to living organisms, it uses an uncalibrated stereo vision system, but no world coordinates and no inverse perspective or kinematic transformations, and it enables an automatic adaptation to changing optical, mechanical and control parameters of the robot.

We implemented the approach on a vision-guided calibration-free articulated arm manipulator with five degrees of freedom. In real-word grasping experiments we showed that robot control without calibration is, indeed, feasible and that it leads to robots with a remarkable degree of robustness against disturbances of the sensors and unexpected changes of the robot's parameters.

Introduction

Conventional vision-based methods for grasping objects typically require a careful calibration of the manipulator and the vision system. Such a calibration tends to be rather cumbersome, and thus, expensive. Moreover, because neither system is perfectly stable, the calibration has to be repeated after relatively short time intervals, and also after any maintenance. The necessity of repeated calibrations is a major impediment to a practical application of vision-guided intelligent robots in industrial and other less structured environments. A robot not requiring any calibration of its optical or mechanical subsystems and able to adapt to changing conditions would be of great practical advantage.

In this paper we describe how such a robot can be realized. In contrast to earlier implementations of our approach we now control all 5 degrees of freedom of the robot and, thus, enable the robot to manipulate a much larger variety of objects of different shapes (Figure 1) and nearly arbitrary initial orientations.

Related Work

Object grasping without calibration has long occupied the attention of researchers in vision-based robot control. Recently a variety of methods has been developed for this challenging and difficult task. For example, [Hollinghurst, Cipolla 1994] and

Figure 1
The gripper of the robot and objects that were used in our experiments, as seen by one of the robot's cameras

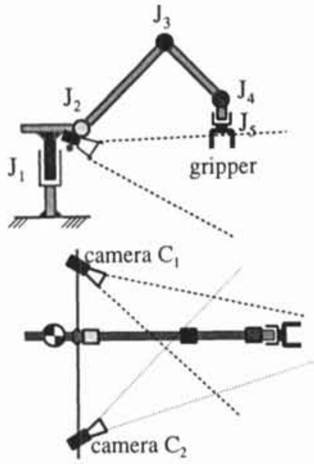

Figure 2
Schematic sketch of the robot arm used for the experiments, its joints and the arrangement of the cameras and their fields of view. The axes of the joints, J_2, J_3 and J_4, are horizontal and parallel to each other. The camera arrangement is not necessarily symmetric.

[Cipolla, Hollinghurst 1997] do not need to use an exact kinematic model of the manipulator or knowledge of the camera parameters because they perform a self-calibration at four known points, combined with the use of visual feedback.

To avoid the necessity of calibrations for robots [Graefe 1995] and [Nguyen, Graefe 1997] have proposed try-and-iterate methods for robust and adaptive robot control. These methods allow the manipulation of objects of various shapes and in any chosen orientation in a horizontal plane. They treat each degree of freedom individually and use qualitative, but not quantitative, knowledge of the robot's kinematics and the camera arrangement.

[Graefe, Maryniak 1998] have implemented the calibration-free control approach underlying those methods in a different way and introduced the Sensor-Control Jacobian matrix to allow a unified control of all degrees of freedom of a robot without using even qualitative knowledge of its kinematics. In their demonstrations they were limited to the control of robots with three or four degrees of freedom only and, therefore, to the grasping of flat cylindrical objects or objects the major axis of which lies in a horizontal plane.

The vision system delivers more information than the minimum necessary for controlling the robot. [Maryniak, Graefe 1998] and [Maryniak, Graefe 1999] present methods for utilizing the redundant information to improve the robustness of the system.

Calibration-Free Grasping

The Problem

A control system for a vision-guided robot is to be designed that does not depend on any quantitative knowledge of the robot's characteristic parameters and, therefore, does not require any calibration. Specifically, in contrast to classical approaches to robot control, no knowledge should be required regarding:

- the exact locations of the cameras
- the exact viewing directions and the internal parameters of the cameras
  (except that both cameras should have the actual work space of the robot in their fields of view)
- the dimensions, kinematics and joint angles of the robot
  (except that, for practical reasons, knowledge that the robot is, e.g., of the articulated arm type, knowledge of the general type of the gripper and knowledge of the number of degrees of freedom may be used if needed)
- the quantitative relationships between the control words sent to the motor controllers and the resulting motions
  (except that these relationships are assumed to be "smooth").

A Solution for Special Cases

If the objects to be grasped are accessible from above and have a vertical axis of rotational symmetry, they may be grasped with the gripper hanging down vertically, while the orientation of the gripper around its axis (joint J_5, Figure 2) is irrelevant. Since in this special case it is sufficient for a successful grasp to make two points, the center point of the open gripper and the reference point of the object, coincide, only three degrees of freedom of the robot have to be controlled. An effective control strategy is then to make the apparent center point of the gripper coincide *in the images of both cameras* with a suitably chosen reference point on the images of the object to be grasped.

Motions of the robot's joints are effected by sending control commands to the motors. The direction and magnitude of the resulting motion depends (in our case, in an unknown way) on a numerical parameter contained in each command. The robot learns the correct values of those parameters automatically and without any previous knowledge by first sending arbitrary control commands to the motors and observing the resulting motions of the gripper in the camera images. For details see [Graefe 1995] and [Graefe, Maryniak 1998].

A More General Case

To grasp objects that lack rotational symmetry, but have such a shape and orientation that they may be grasped from above, a fourth degree of freedom, the one associated with joint J_5, must be activated. To grasp the objects shown in Figure 1 *in all possible orientations*, an arm with 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 orientations.

The method we have developed for this is an extension of the previous one. First, by controlling joints J_1, J_2 and J_3, the gripper is brought to a location (in the images of both cameras) where the image of its center point is near the image of the object reference point. Then the gripper is aligned with the main axis of the object in both images by controlling joints J_4 and J_5. Generally, when two lines in an image are parallel to each other, such as the edges of the gripper and the main axis of the object, it does not imply that the corresponding lines in the world are parallel, too. However, when in the final phases of the grasping process the gripper is close to the object, this error is small enough to be neglected.

Due to the kinematics of the articulated arm, the orientation and location of the gripper cannot be controlled independently of each other. Therefore, the process must be iterated until both the orientation and the location of the gripper are simultaneously approximately correct.

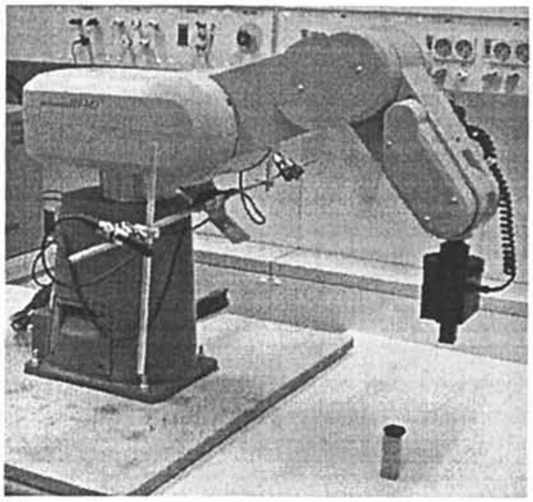

Figure 3
The robot arm used for the experiments. The camera support was deliberately designed to allow arbitrary modifications of the camera arrangement and to make the impossibility of any precise camera adjustment obvious.

The main sensor modality used for the grasping process is stereo vision. The vision system performs the following tasks:

a) Object detection, i.e., separating the gripper and other objects from the background and other irrelevant image information and locating an object that may be grasped

b) Classification of differently shaped objects according to the way how they should be grasped: cylindrical, non-cylindrical, convex, concave etc.

c) Determination of the object's reference point and its orientation in the images of both cameras

d) Identification of the gripper

e) Determination of the gripper center point and the orientation of the gripper in the images

f) Determination of the priority of grasping multiple objects, i.e., the selection of the object to be grasped first when there is more than one object to be grasped

The software of the vision system was realized as a set of subtasks that may run sequentially or in parallel, depending on the computer hardware. It is described in more detail in [Nguyen, Graefe 2000].

Experiments and Results

We evaluated the new approach in real-word experiments where an articulated arm robot with 5 degrees of freedom (Figure 3) was to grasp objects of various shapes as shown in Figure 1. The main sensors of the robot are two video cameras which are mounted on the arm and participate in the rotation of the arm around its vertical axis, but are fixed relative to the work plane of the robot (the plane containing the axes of the joints, J_1 and J_5, Figure 2). All optical, mechanical and control parameters of the system are unknown to the designer and, of course, to the robot.

The images from the two cameras are processed by a vision system, implemented on two processors TMS320C40 (Texas Instruments), one processor for each camera. The vision system detects and tracks relevant objects, determines their locations and orientations in image coordinates and classifies them according to their shape.

The objects that were to be grasped were placed on supports of unknown height (0 to 10 cm) in various almost arbitrary orientations somewhere in the robot's 3-D work space. Due to the lack of a sixth degree of freedom and to certain limitations of the vision system (e.g., slim long objects should not point into one of the cameras), some object orientations had to be avoided.

As expected, in the experiments the objects were searched, detected, located, classified and grasped reliably, regardless of their initial locations and orientations in the robot's work space. To verify the robustness of the robot against unexpected parameter changes two experiments were performed. In the first one the viewing directions of the cameras were arbitrarily modified by large angles, limited only by the requirement, that the gripper and the object to be grasped had to remain in the fields of view of both cameras. In the second experiment the cameras that are normally mounted at equal heights above the table were misadjusted in such a way that one of them was about twice as high above the table as the other one. In both experiments the changes were in no way modeled or communicated to the robot. Nevertheless, the robot continued to function and to grasp objects as before.

Summary and Conclusions

A vision-guided robot that does not use any quantitative model of itself and, therefore, never needs any calibration has been introduced. Due to the lack of a quantitative model and the resulting independence of any calibration, the robot is extremely robust against even major changes in its characteristic parameters. During the actual use of a robot such changes might be caused, e.g., by maintenance operations or the replacement of parts. In experiments it was shown that even major disturbances of the camera orientations could be tolerated whilst the robot was operating. This robustness has the potential for substantial cost savings in the operation of robots in the real world.

The robot is able to manipulate a large variety of objects that may exist anywhere in its work space and have an almost unrestricted orientation. To allow an automatic selection of a suitable grasping strategy for each object, the vision system broadly classifies the objects (cylindrical, non-cylindrical, convex, concave etc.). It then determines suitable grasping points and the orientation of the object in the images of both cameras.

Outlook

Calibration-free robot control will show its full potential when it is used for mobile robots that are to operate in complex and varying environments. Therefore, we hope to implement this approach on our humanoid service robot *HERMES* [Bischoff 2000]. Considering the complexity of *HERMES* (a movable camera head, 2 arms and an omnidirectional wheel base; altogether 25 degrees of freedom), this is a challenging task.

References

Bischoff, R. (2000): Towards the Development of 'Plug-and-Play' Personal Robots. 1st IEEE-RAS International Conference on Humanoid Robots. MIT, Cambridge, MA. CD-ROM.

Cipolla R.; Hollinghurst (1997): Visually Guided Grasping in Unstructured Environments. Robotics and Automation, Elsevier Pub. Pp. 337-346

Graefe, V. (1995): Object- and Behavior-Oriented Stereo Vision for Robust and Adaptive Robot Control. Proc. International Symposium on Microsystems, Intelligent Materials, and Robots, Sendai, pp 560-563.

Graefe, V.; Maryniak, A. (1998): The Sensor-Control Jacobian as a Basic for Controlling Calibration-Free Robots. Proc. IEEE International Symposium on Industrial Electronics, ISIE'98, Pretoria, South Africa, pp. 420-426.

Hollinghurst, N. J.; Cipolla, R. (1994): Uncalibrated Stereo Hand-Eye Coordination. Image and vision Computing 12 (3), pp. 187-192.

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; Proceeding of the CWCICIA'97, Vol. 1, Xian-China, pp. 77-83.

Nguyen, M.-C.; Graefe, V. (2000): Stereo Vision-Based Object Classification. Applications. IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS '00. Takamatsu.