

# Integration of Vision and Force Sensors for Grasping

Peter K. Allen   Andrew T. Miller   Paul Y. Oh   Brian S. Leibowitz

Department of Computer Science, Columbia University, New York, NY 10027 \*

## Abstract

This paper describes a set of methods that can be used to integrate real-time external vision sensing with internal force and position sensing to estimate contact forces by the fingers of a hand. Estimating these forces and contacts is essential to performing dextrous manipulation tasks. Most robotic hands are either sensorless or lack the ability to accurately and robustly report position and force information relating to contact. By adding external vision sensing, we can complement any internal sensors to more accurately estimate forces and contact positions. Experiments are described that use real-time visual trackers in conjunction with internal strain gauges and a new tactile sensor to accurately estimate finger contacts and applied forces for a three fingered robotic hand.

## 1 Introduction

Robotic dextrous manipulation remains a difficult problem with many open research problems. While there have been a number of detailed analyses of the kinematic and dynamic constraints necessary to effect stable grasps, most require a high level of sensory input and feedback from the grasping device (i.e. robotic hand) to perform dextrous manipulation. The sensory information required typically includes contact point estimation, surface normal and curvature measures, and knowledge of both applied and induced forces on the fingers of the hand. Any hand system that is capable of providing this information is by definition a multi-sensor device, and requires a set of software and hardware integration modules to properly estimate these measures. While great strides have been made in robotic hand design and a number of working dextrous robotic hands built, the reality is that the sensory information required for dextrous manipulation lags the mechanical capability of the hands. Accurate and high bandwidth force and position information for a multiple finger hand is still difficult to acquire robustly.

---

\*This work was supported in part by DARPA contract DACA-76-92-C-0007, NSF grants CDA-90-24735 and IRI-93-11877, and an ONR MURI Grant

This paper describes a set of methods that integrate vision, force, and tactile information for grasping tasks. The high sensory demands required by grasping motivate the use of a multi-sensor fusion approach. This is especially true since robustness is a key issue in grasping. Our aim is to use a set of internal and external sensors to estimate forces and contacts on the hand. Vision can be an effective sensing modality for grasping tasks due to its speed, low cost, and flexibility. It can serve as an external sensor that can provide control information for devices that lack internal sensing or that would require extensive modification and re-engineering to provide contact and force sensing. Without this sensory feedback, open loop control must be used which requires precise models of the environment to be effective. Using a vision system, a simple uninstrumented gripper/hand can become a precision device capable of position and possibly even force control. Additionally, when vision is coupled with any existing internal hand sensing, it can provide a rich set of complementary information to confirm and quantify internal sensory data.

## 2 Related Research

A number of previous researchers have explored the use of visual feedback and control to assist in the grasping task. Houshangi [7] tracked moving objects for grasping. Hollingshurst and Cipolla [6] have developed a system for positioning a gripper above an object in the environment using an affine stereo transform to estimate the object's position. Taylor et al. have used 3-D vision to guide the grasping task [14]. Castano and Hutchinson [3] use visual constraint planes to create compliant surfaces for constraint robot movement in the real world. Bendiksen and Hager [2] have used vision in conjunction with gripper forces to achieve stable grasps. Sharma et al. [10] use perceptual 3D surfaces to represent the workspace of the gripper and object and they plan their positioning tasks along these surfaces. Sobh and Bajcsy [12] examined how finite state machines can be used to monitor the grasping process through vision. Smith and Papanikolopoulos [11] have recently extended their visual servoing and control algorithms to create a hand-eye tracker capable of grasping static and moving ob-

jects. Our own work has explored the capability of vision systems to track and grasp moving objects [1] and use uncalibrated visual servoing to perform alignment tasks [16]. This work motivated us to use stereo vision to control an uninstrumented gripper in simple grasping tasks [17, 15]. The robotic hand used in that research had no internal sensing, and thus vision system could not report forces being applied to the grasped objects. In this paper, we explore the integration of external vision with existing but limited internal force and contact sensors.

### 3 Kinematic Issues

Integration of position information from vision sensors for grasping requires an accurate kinematic model of the hand. The kinematic model is used to support and interpret force-related sensory information that is the result of fusing visual contact information with internal force sensing. Any implementation of grasping for a dextrous robot hand demands knowledge of both the internal and external forces and torques acting on the hand. Below, we describe this kinematic and force model for a particular robotic hand, but the method can be extended to other hands as well.

#### 3.1 The Barrett Hand

The dextrous robot hand used for our investigation is the Barrett Hand depicted in Figure 1. It is an eight-axis, three-fingered mechanical hand with each finger having two joints. One finger is stationary and the other two can spread synchronously up to 180 degrees about the palm (finger 3 is stationary and fingers 1 and 2 rotate about the palm). Although there are 8 axes, the hand is controlled by four motors. Each of the three fingers has one actuated “inner” link, and a coupled “outer” link that moves at a fixed rate with the inner link. A novel clutch mechanism allows the outer link to continue to move if the inner link’s motion is obstructed (referred to as breakaway). An additional motor controls the synchronous spread of the two fingers about the palm. Various grasp classifications capable with the hand include but are not limited to: power, hook, capture, cylinder-tip, spherical, and cylinder grasps.

#### 3.2 Internal Force Sensing

Each finger has been equipped with a strain gauge. The gauges measure the axial strain in response to loads applied to the outer link due to coupled cable tensions (see Figure 2). The beam to which the strain gauges are attached to can be modeled as a cantilever beam under static equilibrium as follows.

$$-T_1 \sin \gamma - T_1 \sin \gamma + T_2 \sin \gamma + T_2 \sin \gamma - F_{FT} = 0 \quad (1)$$




Figure 1: The Barrett Hand




Figure 2: Location of Strain Gauges and Tensions

$$F_{FT} = 2(T_2 - T_1) \sin \gamma = 2\Delta T \sin \gamma \quad (2)$$

where  $F_{FT}$  is the force along the finger’s outer link,  $T_1$  and  $T_2$  are the tensions along the upper and lower cables respectively, and  $\gamma$  is the directional angle these tensions act along. Thus, given a strain gauge reading  $\Delta T$ ,  $F_{FT}$  can be calculated.

For point forces acting in the x-y plane, the General Beam Method can be used to determine the deflection at any point along the finger’s outer link given  $F_{FT}$  and the location of the applied force. This is illustrated in Figure 3. Recall that beam curvature is related to a bending moment  $M_b$  as

$$\frac{d^2y}{dx^2} = \frac{M_b}{EI_{zz}} \quad (3)$$

where  $E$  is Young’s modulus and  $I_{zz}$  is the moment of inertia about the z-axis. Given clamped boundary conditions, the deflection at any point along the beam is given by

$$y(x) = \frac{F_{FT}x^3}{6EI_{zz}} - \frac{F_{FT}Lx^2}{2EI_{zz}} \quad (4)$$

where  $L$  is the length of the beam. Maximum deflection occurs at  $x = L$ , that is, at the finger tip:




Figure 3: Free Body Diagram

$$y(L) = -\frac{F_{FT}L^3}{3EI_{zz}} \quad (5)$$

### 3.3 Predicting Forces from Strain Gauges

It is desired to solve for the applied forces on the finger,  $F_{FT}$  as introduced in (4) without explicitly measuring the vertical finger displacement  $y(x)$ . This can be done by assuming that the strain gauge output  $s_g$  and  $y(x)$  are linear such that

$$s_g = k_1 F_{FT} x^3 - k_2 F_{FT} x^2 + a + \varepsilon \quad (6)$$

where  $a$  is an offset and error  $\varepsilon$ . Formulating a least-squares parameter estimate  $\hat{\Theta}$  where

$$\Phi = \begin{bmatrix} F_{FT}[1]x[1]^3 & -F_{FT}[1]x[1]^2 & 1 \\ F_{FT}[2]x[2]^3 & -F_{FT}[2]x[2]^2 & 1 \\ \vdots & \vdots & \vdots \\ F_{FT}[n]x[n]^3 & -F_{FT}[n]x[n]^2 & 1 \end{bmatrix} \quad (7)$$

$$\hat{\Theta} = [k_1 \quad k_2 \quad a]^T \quad (8)$$

$$y = [s_g[1] \quad s_g[2] \quad \dots \quad s_g[n]]^T \quad (9)$$

such that

$$\hat{\Theta} = (\Phi^T \Phi)^{-1} \Phi^T y \quad (10)$$

The functional matrix  $\Phi$  is formed from values  $F_{FT}$  and  $x$  used in calibrating the strain gauges. In this calibration process, fixed weights ranging from 1.2 lbs (4.72 N) to 4 lbs (18.09 N), at 0.2 lbs intervals, were suspended along the finger at 2 mm intervals.

Since the strain gauges as illustrated in Figure 2 are configured in a conventional Wheatstone bridge, it is reasonable to assume linearity between  $s_g$  and  $y(x)$ . Figure 4 is a plot of the finger force vs. strain gauge and distance function. By kinematically deriving the predicted




Figure 4: Finger Force as a function of strain gauge readings ( $s_g$ ) and distance ( $x$ )

beam behavior, and empirically determining the parameter estimates, a predictive force model can be solved for without explicitly measuring the vertical displacement as:

$$F_{FT} = \frac{s_g - a}{k_1 x^3 - k_2 x^2} \quad (11)$$

Furthermore the moment of inertia  $I_{zz}$  for the finger does not need to be explicitly calculated. This would pose to be challenging given both the non-uniform cross sectional area and non-homogeneous composition of the finger. Experimental data in Section 4.2 demonstrates the validity of the above predictive force model.

## 4 Integration of Vision with Force

The strain gauge output can report the forces acting on the outer link, but cannot localize them. The strain gauges only provide us with torque readings about the outer joints - it is necessary to find the point of contact along a finger to determine the force normal to that finger. Vision sensing can be used to provide this contact point estimation, and thereby calculate the actual finger tip forces in conjunction with the strain gauge readings and the kinematic model developed in the previous section. The effective requirements of any vision sensing system to perform this integration task include real-time contact position estimation of multiple contacts. Since the fingers are moving until contact is made, tracking the fingers and objects to be grasped as they move over time is also required.

### 4.1 Vision Sensor

The initial vision sensor we are using is a single camera that can image the fingers of the hand and the object to be contacted. In this paper, we are using a scaled orthography camera model which effectively allows us to determine the 3-D position of fingers and contacts from the image plane directly. The methods described here

can be easily extended to a stereo perspective camera model that can provide full 3-D depth recovery as well [15]. The vision system we are using is a modification of Hager's X Vision system [5]. X Vision was chosen because it is a software implementation that can be easily applied to an existing hand system simply by hooking up a camera and frame grabber. It also has the advantage of portability across a number of platforms. Using X Vision, complex trackers can be built up from simple edge and region-based tracking primitives. Each tracker has a state vector consisting of position and orientation information which is updated after each iteration of the tracking loop. Once a line or region tracker is initialized on an edge or window within the image it will track the feature at frame rates. The tracker assumes an affine transformation (translation, rotation, scaling and shear) between successive image frames.

These trackers have several parameters which can be altered by the user. The length of the line tracker, as well as the width of the local window around the line, can be set. In the region based SSD (Sum of Squared Distances) correlation tracker many parameters affect its performance. These include window size and resolution, as well as several flags to enable or disable rotation, scaling, and brightness compensation. By seeding these trackers to follow edges or regions in the image, they will continuously track the features in real time. The trackers are very reliable for reasonable amounts of motion. In our experiments we used three different trackers: a corner tracker that is composed of two line trackers that reflect the intersections of the lines along the inner and outer links of a finger, another line tracker to track a spike-like force probe that contacts the finger, and an SSD tracker for the endpoint of the finger. Using the scaled orthographic camera model, we were able to track a finger and determine a point of contact to within 1 mm of the actual contact point.

## 4.2 Experiment I

The purpose of this experiment was to see if vision sensing could track finger and object movements, and then localize contact along the finger to estimate actual finger forces using the calibration data from the previous section. We mounted a spike-like force probe on top of an ATI force sensing wrist that provided us with accurate three dimensional force data which we used as ground truth. The hand was mounted on a PUMA-560 robot and positioned in the vicinity of the spike. One finger of the hand was positioned above the spike as the trackers were initialized. To find the point where the spike contacts the finger we used one line tracker initialized on the right side of the spike, one corner tracker placed along inside edges of the finger, and an SSD tracker initialized so that it is centered on a point marked at 7 mm from the end of the finger which is used as a reference (see




Figure 5: Finger of Barrett Hand applying force to instrumented probe




Figure 6: Visual Sensing: Line tracker, corner tracker, and SSD tracker applied to finger

Figures 5 and 6).

As the finger closed on top of the spike, the trackers followed it. The corner tracker's state vector reported the  $(x, y)$  position of the intersection of its two lines as well as their individual orientations with respect to the  $x$ -axis. Although the orientation of the line tracking the inner link was not needed for this experiment, we have found that the corner tracker is more robust than a simple line tracker because the endpoint intersection constraint prevents each of the lines from sliding along their respective edges. In later experiments the two lines of the corner tracker can be used to compute the angle of the outer joint simply by finding the difference in orientation between the two lines.

Using the state vectors of the trackers (see Figure 7), a point of intersection in image space was computed:

$$y = L1.y - offset, \quad x = \frac{y - C.y}{\tan(C.\theta_2)} + C.x \quad (12)$$

where  $L1.y$  is  $y$ -coordinate of the line tracker follow-




Figure 7: Using three feature trackers, the point of contact can be computed.

ing the spike,  $(C.x, C.y)$  is the position of corner tracker,  $C.\theta_2$  is the orientation of the upper line in the corner tracker, and *offset* is the distance in pixels from the center of the spike to the edge. The point of finger contact and the point tracked by the SSD were then transformed to world coordinates using the scaled orthography camera calibration matrix, and a distance along the finger was computed.

This method resulted in fast, reliable measurements within 1mm of the actual point of contact. Near the end of the finger where the distance between the contact point and the tracked endpoint reference was small, our accuracy dropped to about 1.5mm from 1mm. Using 10 averaged distance measurements and 10 averaged strain gauge readings taken from the hand we were able to predict the force applied to the finger at the point of contact. The standard deviation in position was .4 mm along the finger except at the very tip where it increased to .9 mm. By varying the velocity of movement of the finger, we were able to apply a range of forces on the spike, limited by maximum torque available from the motor. Our accuracy was also limited by the low resolution of the strain gauges in this force range. The maximum range was near four bits at the end of the finger where the maximum deflection of the beam occurs and results in a corresponding higher strain gauge reading. Contacts near the joint between the outer and inner links caused the least deflection and the strain gauge readings were less accurate in this region, only about two bits of information.

Our results (see Figure 8) showed that as we placed the spike further along the finger our error dropped until we reached the very end of the finger. The low resolution of the strain gauges when the force acts near the joint also makes calibration in this region difficult. Generally we found that as more pressure is applied to the spike, our error percentage dropped significantly.




Figure 8: Results of Experiment I: Graphs (top to bottom) show force data when spike is at 13mm from outer joint, 24mm, 40mm, and 51mm. Solid Line represents modeled linear relationship between strain gauge values and forces for a contact at the given distance along the finger. Points are actual force readings from the ATI force sensing wrist for contact at the distance determined by vision. Contacts farther out on the link were modeled more accurately due to the limited range of the finger forces and strain gauge resolution near the finger joint.




Figure 9: Tactile sensor mounted on Barrett Hand.

### 4.3 Experiment II

While the previous experiment showed the utility of vision in tracking fingers and determining contact points, it is important to note that occlusion can prevent the vision sensor from reporting this information. This can be overcome by the use of a finger mounted tactile sensor that can estimate contact point localization. We are using a capacitive tactile sensor designed by Howe [13] that is based upon an earlier design of Fearing's [4]. The sensor is designed to be slipped on to the outer link of the finger as shown in Figure 9. We are currently building additional sensors for the inner links and palmar surfaces of the hand. The tactile sensor geometry is a 4x8 grid with each capacitive cell approximately 3 mm by 3 mm and 1 mm spacing between tactile elements (tactels), and the sensor can bend to the curve of the fingertip. The sensor is covered with a compliant elastomer that allows force distributions to be spread over the sensing surface. While the intensity values have not yet been calibrated to provide accurate force information, it is possible to use their relative responses to compute the weighted center of contact of the applied force using moments. To calibrate the sensor to provide the point of contact, we probed one column of the sensor which was aligned along the length of the finger at different distances along its length and it reported the contact center with a precision of about 1.2 mm. Figure 10 shows the predicted relationship between tactel location and distance along the finger and sample probes along the sensor attached to the outer link.

Tactile sensors can also be used without vision to estimate forces and contacts [8, 9]. However, the robustness of these sensors is always a problem, and the vision sensing can serve to verify as well as improve the tactile localization, particularly in dead spots between tactels.




Figure 10: Contact distance vs. tactile localization along a single column of the tactile sensor.

### 4.4 Experiment III

One of the Barrett Hand's unique features is its ability to complete a grasp of an object after the inner link's motion has been obstructed. During this breakaway situation, the clutch mechanism in each finger allows the outer links to continue closing after its inner link has been stopped. This feature is especially useful for grasping large objects or irregularly shaped objects where the inner links may be blocked early, and the outer links finish the grasp.

Each of the four motors in the hand is equipped with an optical position encoder that can supply angular data for the finger joint where it meets the palm (referred to as the inner joint). Using the kinematic parameters of the hand, this allows us to calculate the endpoint of the inner link. There is no encoder at the joint between the inner and outer links (referred to as the outer joint). In normal operation, the outer joint of each finger is driven at a 4:3 ratio with respect to the inner joint, and using the kinematic equations and this angle ratio, the finger's position in space can be computed. Often, however, due to backlash, this information is not exact. Moreover, in a breakaway situation, the angle of the outer joint cannot be derived from the optical encoder in the inner joint since the clutch has disengaged the links and the 4:3 ratio no longer is valid. Vision sensing can provide the joint angle by calculating the difference in orientation angle for the two line trackers along the inner and outer links of the finger. Knowing this joint angle and the link geometry allows us to use forward kinematics and locate the last link in space. During breakaway, without vision sensing, we cannot localize the contact along the outer link.

In this experiment we rigidly mounted a small block to the palm of the hand and closed the third finger around it. By securing the block from sliding, we were able to better ensure the obstruction of the first link, and cause the grasp to result in two contacts by the inner and outer links on the block which we attempted to locate. The motion trackers in this experiment included

two SSD trackers and three corner trackers. One SSD tracker was initialized on the end of the finger and a second one was centered on a cross hair marking the inner joint's axis. Two corner trackers were placed on the block, one on each side, and the third corner tracker was placed along the inside edges of the finger as before. After the trackers had been initialized, the finger was closed around the block (see Figures 11, 12, and 13). Using this method, we were able to track the angle of the joint between inner and outer links at all times. At this point the points of contact could be computed by using a method similar to the previous experiment. First the intersection of each side of the block with the two lines of the corner tracker was computed and then using the SSD trackers as reference points, the distance along the link could be determined. We found that the vision system reported contact points that were within 2 mm of the actual contact points. We encountered some difficulty in keeping the first SSD tracker centered on the mark near the end of the finger. In future revisions of this system we hope to improve the tracker or the pattern being tracked and thus improve our accuracy. We also attached the tactile sensor to the outer link of finger 3 to provide us with additional contact information. In this experiment, the actual contact distance along the outer link was determined to be 38 mm as read by a ruler. The vision system reported 40.0mm as the distance and the tactile sensor reported it as 36.6mm. The actual contact distance along the inner link was determined to be 49 mm as read by a ruler and the vision system reported 51.3 mm as the distance (we did not have an inner link tactile sensor).

## 5 Conclusions and Future Work

As stated above, grasping with a robotic hand is an inherent multi-sensor fusion problem. Humans appear to have very easily integrated position, force and vision sensing during grasping. One way to extend these capabilities to robotic hand systems is to add complementary sensors that provide both unique and redundant information relating to force and position. In this paper, we have described a set of visual trackers that can be used in a control loop to 1) determine if contact has occurred and 2) estimate the position of contact along a finger. Vision has the advantage of being fast and simple to add to an existing hand system that may be lacking in sensory capability. We have also calibrated the internal force sensors of the hand to fuse the visual contact position information with the modeled strain gauge values and correctly predict applied finger forces. This will allow a grasping task to 1) visually determine the grasping points of contact and 2) modify the control of the finger to apply pre-determined forces on an object. Finally, we have also integrated a tactile sensor on the outer link




Figure 11: The trackers after initialization.




Figure 12: The inner link is stopped, and breakaway occurs.




Figure 13: The outer link completes the grasp and the contact points are computed.

of the finger to complement the position sensing of the vision system. The tactile sensor can report positional contact location also (in a limited region of coverage), which can be used when vision is occluded, as well as confirming any visually determined contact positions.

We are currently extending the system described here in a number of ways. First, we are extending the vision system to a 3-D binocular stereo setup. By using pairs of feature trackers, one in each image, we will be able to compute depth information in the scene. Also, by using more than one camera for external visual sensing, we can reduce the occlusion problems during grasping. We are also extending these visual trackers to track multiple finger features robustly. Second, we are applying these methods to a set of grasping tasks that include handling deformable and fragile objects, where finger forces need to be pre-determined and monitored. Third, we are instrumenting all three fingers of the hand with tactile sensors on inner and outer links, as well as a palmar sensor. We still need to characterize the finger forces from the tactile responses. Our previous work with tactile sensors has shown that this is often difficult to do accurately and robustly; we hope the multi-sensor approach described above will allow us to complement the tactile responses with strain gauge and vision sensing for robustness. This will also include investigating the role of dynamics of the fingers and sensors for integrating force and vision.

**Acknowledgment:** We would like to thank Rob Howe, Jae Son, Bill Peine, and Parris Wellman from the Harvard Robotics Lab for their help in building the tactile sensor system.

## References

- [1] P. Allen, A. Timcenko, B. Yoshimi, and P. Michelman. Automated tracking and grasping of a moving object with a robotic hand-eye system. *IEEE Trans. on Robotics and Automation*, 9(2):152-165, 1993.
- [2] A. Bendiksen and G. Hager. A vision-based grasping system for unfamiliar planar objects. In *IEEE International Conference of Robotics and Automation*, pages 2844-2849, May 1994.
- [3] A. Castano and S. Hutchinson. Visual compliance: Task-directed visual servo control. *IEEE Trans. on Robotics and Automation*, 10(3):334-342, June 1994.
- [4] R. S. Fearing. Tactile sensing for shape interpretation. In S. T. Venkataraman and T. Iberall, editors, *Dextrous Robot Hands*. Springer-Verlag, 1989.
- [5] G. D. Hager and K. Toyama. X vision: A portable substrate for real-time vision applications. Technical report, Department of Computer Science, Yale University, 1995.
- [6] N. Hollinghurst and R. Cipolla. Uncalibrated stereo hand-eye coordination. Technical Report CUED/F-INFENG/TR126, Department of Engineering, University of Cambridge, 1993.
- [7] N. Houshangi. Control of a robot manipulator to grasp a moving target using vision. In *IEEE International Conference on Robotics and Automation*, pages 604-609, Cincinnati, May 13-18 1990.
- [8] D. Johnston, P. Zhang, J. Hollerbach, and S. Jacobsen. A full tactile sensing suite for dextrous robot hands and use in contact force control. In *Proc. of the 1996 IEEE International Conference on Robotics and Automation*, pages 3222-3227, 1996.
- [9] H. Maekawa, K. Tanie, and K. Komoriya. Dynamic grasping force control using tactile feedback for grasp of multifingered hand. In *Proc. of the 1996 IEEE International Conference on Robotics and Automation*, pages 2462-2469, 1996.
- [10] R. Sharma, J. Herve, and P. Cucka. Analysis of dynamic hand positioning tasks using visual feedback. Technical Report CAR-TR-574, Center for Auto. Res., University of Maryland, 1991.
- [11] C. Smith and N. Papanikolopoulos. Vision-guided robotic grasping: Issues and experiments. In *Proc. of the 1996 IEEE International Conference on Robotics and Automation*, pages 3203-3208, 1996.
- [12] T. M. Sobh and R. Bajcsy. Autonomous observation under uncertainty. In *Proc. of IEEE International Conference on Robotics and Automation*, pages 1792-1798, May 1992.
- [13] J. Son and R. Howe. Tactile sensing and stiffness control with multifingered hands. In *Proc. of the 1996 IEEE International Conference on Robotics and Automation*, pages 3228-3233, 1996.
- [14] M. Taylor, A. Blake, and A. Cox. Visually guided grasping in 3d. In *IEEE International Conference of Robotics and Automation*, pages 761-766, May 1994.
- [15] B. Yoshimi. *Visual Control of Robotics Tasks*. PhD thesis, Dept. of Computer Science, Columbia University, 1995.
- [16] B. Yoshimi and P. Allen. Alignment using an uncalibrated camera system. *IEEE Trans. on Robotics and Automation*, 11(4):516-521, August 1995.
- [17] B. H. Yoshimi and P. Allen. Visual control of grasping and manipulation tasks. In *MFI '94: 1994 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, pages 575-582, 1994.